Top 10 Free Windows Server Administration Tools

Windows IT Pro

A PENTON PUBLICATION

APRIL 2013 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

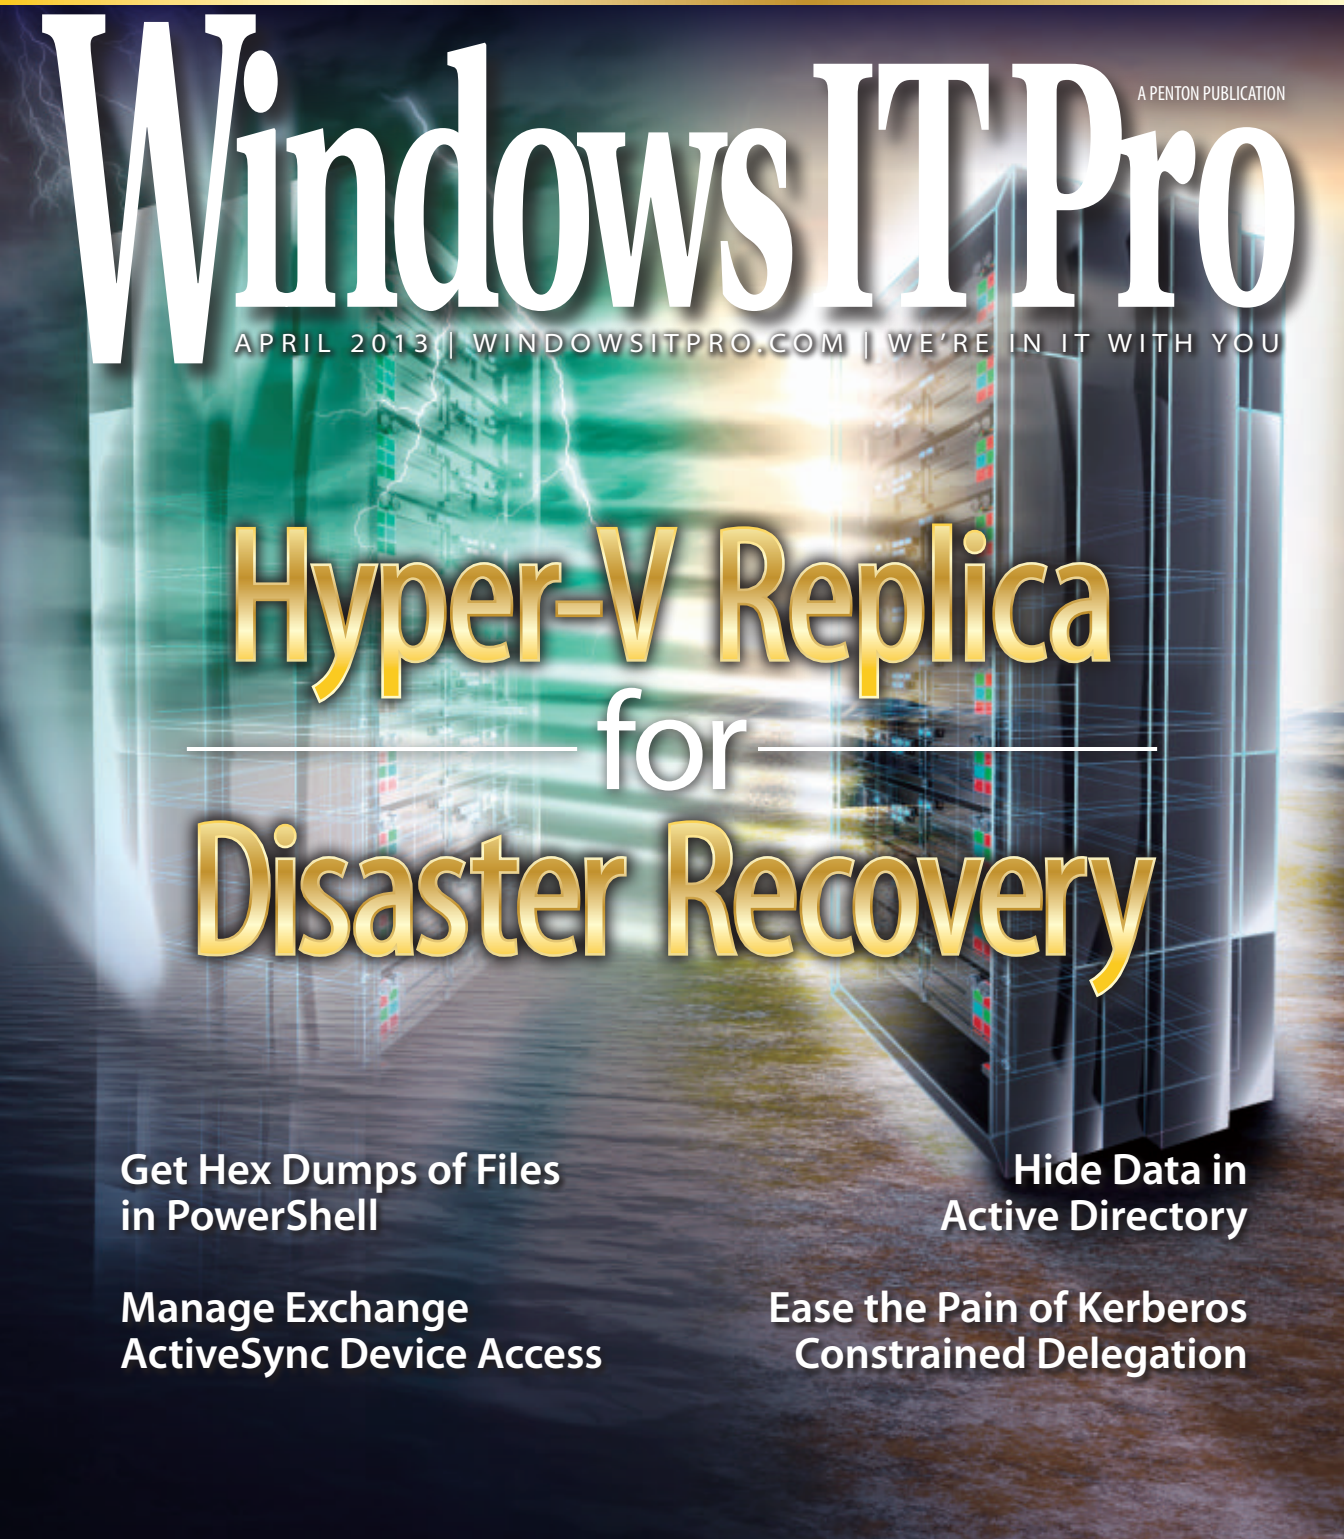

Hyper-V Replica for Disaster Recovery

Get Hex Dumps of Files in PowerShell

Hide Data in Active Directory

Manage Exchange ActiveSync Device Access

Ease the Pain of Kerberos Constrained Delegation

1&1 DYNAMIC CLOUD SERVER

SAVE $120 First year*

1&1 DYNAMIC CLOUD SERVER

A fully flexible server for a range of requirements including applications, databases, gaming and much more!

- Independently configure CPU, RAM, and storage
- Accurate and fair: Control costs with pay-per-configuration and hourly billing
- Up to 6 Cores, 24 GB RAM, 800 GB storage
- 2000 GB of traffic included free
- Parallels® Plesk Panel 11 for unlimited domains; reseller ready
- Up to 99 virtual machines with different configurations under one account
- No setup fee
- 24/7 phone and e-mail support

$39.99 per month* ~~$49.99~~ per month*

MAXIMUM FLEXIBILITY
Change hardware configurations in real time to meet your business needs

MAXIMUM SECURITY
Redundant storage and mirrored processing units reliably protect your server

INCLUDES 1 DEDICATED SSL CERTIFICATE

HOURLY BILLING
Control Costs by only paying for what you configure

PARALLELS PLESK® PANEL 11
for unlimited domains

Parallels Plesk Panel

OPTERON PROCESSOR AMD

SUSE

MOBILE APPS
for server management and monitoring

1&1®

MEMBER OF united internet

1and1.com

*Offer valid for a limited time only. Base configuration includes 1 processor core, 1 GB RAM, 100 GB storage. This offer applies to new contracts only. 12 month minimum contract term. Other terms and conditions may apply. Visit www.1and1.com for full promotional offer details. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 Internet, all other trademarks are the property of their respective owners. © 2013 1&1 Internet. All rights reserved.

CAN'T GET AWAY?

Get first-class education from your desk

Windows IT Pro offers FREE online events including webcasts, demos and virtual conferences. All events are brought to your computer live while being fully interactive.

Go to www.windowsitpro.com/events to see an up-to-date list of all online events.

Windows IT Pro

First-class education from the top experts in the industry.

Visit www.windowsitpro.com/events for a knowledge upgrade today!

Have a full plate on the live date? Don't sweat it! All online events are recorded and available 24/7.

APRIL 2013 / VOL. 19 / NO. 4

COVER STORY

Using Windows Server 2012 Hyper-V Replica 32

— John Savill

Windows Server 2012 Hyper-V Replica lets you asynchronously replicate a virtual machine to a second Hyper-V host—even across geographically separate locations. Find out more about this new feature.

Features

45 Get Hex Dumps of Files in PowerShell
Bill Stewart

51 Managing Exchange ActiveSync Device Access
Paul Robichaux

61 Using the Confidentiality Bit to Hide Data in Active Directory
Guido Grillenmeier

Products

77 New & Improved

Interact

73 Ask the Experts

In Every Issue

81 Ctrl+Alt+Del

83 Advertiser Directory

83 Directory of Services

83 Vendor Directory

Chat with Us

Facebook Twitter LinkedIn

Columns

6 Need to Know

Office 365 for Businesses
Paul Thurrott

12 Windows Power Tools

Get-Credential Gives You Power Over Your Passwords
Mark Minasi

16 Top 10

Free Windows Server Administration Tools
Michael Otey

19 Enterprise Identity

7 More Reasons to Upgrade to Windows Server 2012 Active Directory
Sean Deuby

24 What Would Microsoft Support Do?

How Windows Server 2012 Eases the Pain of Kerberos Constrained Delegation, Part 2
Mike Stephens

Editorial
Editorial Director: Megan Keller
Editor-in-Chief: Amy Eisenberg
Senior Technical Director: Michael Otey
Technical Director: Sean Deuby
Senior Technical Analyst: Paul Thurrott
Custom Group Editorial Director:
Dave Bernard
Exchange & Outlook, Cloud:
Brian Winstead
Systems Management, Networking,
Hardware: Jason Bovberg
Scripting: Blair Greenwood
SharePoint, Active Directory, Security,
Virtualization: Caroline Marwitz
SQL Server, Developer Content:
Megan Keller
Managing Editor: Lavon Peters
Editorial SEO Specialist: Jayleen Heft

Senior Contributing Editors
David Chernicoff, Mark Minasi,
Tony Redmond, Paul Robichaux,
Mark Russinovich, John Savill

Contributing Editors
Alex K. Angelopoulos, Michael Dragone,
Jeff Fellinge, Brett Hill, Dan Holme,
Darren Mar-Elia, Eric B. Rux,
William Sheldon, Curt Spanburgh,
Bill Stewart, Orin Thomas, Douglas Toombs,
Ethan Wilansky

Art & Production
Senior Graphic Designer: Matt Wiebe
Director of Production: Dylan Goodwin
Group Production Manager:
Julie Jantzer-Ward
Project Manager: Adriane Wineinger
Graphic Specialists: Karly Prickett &
Ashley Lawson

Advertising Sales
Technology Market Leader: Peg Miller
Key Account Director:
Chrissy Ferraro • 970-203-2883
Account Executives:
Megan Key • 970-203-2844
Barbara Ritter • 858-367-8058
Cass Schulz • 858-357-7649

Client Services
Senior Client Services Manager:
Michelle Andrews • 970-613-4964
Ad Production Coordinator: Kara Walby

Marketing & Circulation
Customer Service
Senior Director, Marketing Analytics:
Tricia Syed

Technology Division
& Penton Marketing Services
Senior Vice President: Sanjay Mutha

Corporate
Chief Executive Officer:
David Kieselstein
Chief Financial Officer/Executive Vice
President: Nicola Allais

Penton

List Rentals
MeritDirect
333 Westchester Avenue,
White Plains, NY 10604

Reprints
Reprint Sales:
Wright's Media • 877-652-5295

Windows IT Pro, April 2013, Issue No. 224,
ISSN 1552-3136. *Windows IT Pro* is published monthly by Penton Media, Inc. Copyright ©2013 Penton Media, Inc. All rights reserved. No part of this publication may be reproduced or distributed in any way without the written consent of Penton Media, Inc.

Windows IT Pro, 748 Whalers Way, Fort Collins, CO 80525, 800-621-1544 or 970-663-4700. Customer Service: 800-793-5697.

We welcome your comments and suggestions about the content of *Windows IT Pro*. We reserve the right to edit all submissions. Letters should include your name and address. Please direct all letters to letters@windowsitpro.com. IT pros interested in writing for *Windows IT Pro* can submit articles to articles@windowsitpro.com.

Program Code: Unless otherwise noted, all programming code in this issue is ©2013, Penton Media, Inc., all rights reserved. These programs may not be reproduced or distributed in any form without permission in writing from the publisher. It is the reader's responsibility to ensure procedures and techniques used from this publication are accurate and appropriate for the user's installation. No warranty is implied or expressed.

Windows®, Windows Vista®, and Windows Server® are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries and are used by Penton Media, Inc., under license from owner. *Windows IT Pro* is an independent publication not affiliated with Microsoft Corporation. Microsoft Corporation is not responsible in any way for the editorial policy or other contents of the publication.

Windows IT Pro

Office 365 for Businesses

Paul Thurrott

is senior technical analyst for Windows IT Pro. He writes the SuperSite for Windows, a weekly editorial for Windows IT Pro UPDATE, and a daily Windows news and information newsletter called WinInfo Daily UPDATE.

Email

Twitter

Website

In January of 2013, Microsoft launched Office 365 Home Premium, a new version of Office 365 aimed at households and individuals. Now the firm is updating its business-oriented Office 365 services, all of which offer the power and flexibility of Exchange, SharePoint, and Lync in the cloud, and providing a new low-cost entry for small businesses.

Before I go further into these latest Office 365 updates, there are a few key takeaways for business users curious about how the home offering of Office 365 differs from the business-oriented Office 365 versions. First, the Home Premium subscription is licensed for a household (or individual) and comes with five installs of Office 2013 Professional that can be shared by multiple people in a family. Second, Home Premium doesn't include Exchange Online, SharePoint Online, or Lync Online like the business-oriented versions of Office 365; instead, users are expected to use consumer-oriented services such as Outlook.com and SkyDrive. And third, Office 365 Home Premium is really inexpensive at just $99.99 per year, assuming, of course, you need multiple copies of Office.

Most IT pros are probably familiar with the fact that Microsoft first launched Office 365, solely for businesses of various sizes, a few years back. And most are probably familiar with the basic functionality of these offerings, and the value proposition. So let's focus on what's changed.

Microsoft doesn't provide version numbers for Office 365, part of the company's whole reimagining of its product offerings as online services. But when you consider the changes to both the services themselves and to the Office suite that comes with many of the subscriptions, it's fair to say that this is a major update—a version 2.0 of sorts.

A New Office

Office 365 is being upgraded with the latest Office, Office 2013, or what Microsoft curiously calls "the New Office." This is a major upgrade of

the suite, with amazing new Click-to-Run deployment and installment technologies, a related Office On Demand feature that lets you run the core Office applications "on the fly" from any PC by streaming it temporarily over the Internet, and a crisp, clean new UI that's as at home on traditional PCs as it is on new Windows 8 devices. Office 2013 availability varies between Office 365 subscriptions, but most of the offerings also include a nicely revamped version of the Office Web Apps, web-based versions of Word, Excel, PowerPoint, and OneNote that work well in any modern, HTML5-compliant browser and offer an increasingly sophisticated set of functionality related to document editing and multi-user collaboration.

For the Office 365 subscriptions, those who do qualify for an Office suite will now receive Office 2013 Professional Plus. (Those who don't can license Office Pro Plus for $144.99 per user per year.) This includes Word, PowerPoint, Excel, Outlook, OneNote, Access, Publisher, and Lync 2013. Aside from the previously noted Click-to-Run capabilities, the big change from previous Office 365 versions with Office 2010 is that licensing is now per user, not per device. So each user with access to Office 2013 Pro Plus can install it on up to five devices. Today, that means Windows PCs and devices as well as Macs, but there are strong hints that other devices—such as Apple iPads and possibly Android tablets—are coming as well.

Additionally, the App-V–based Office 2013 installation means that you can run Office 2013 Pro Plus side by side with your previous Office version, which can be helpful for users with finicky add-ins or those who wish to move forward at their own pace.

New Office 365 Subscriptions

In addition to the Office 365 subscriptions that existed previously, and the consumer-oriented Office 365 Home Premium offering that Microsoft announced last month, there are two major new Office 365 versions being offered to businesses this year: Office 365 Small Business Premium and Office 365 for Midsized Businesses.

Office 365 Small Business Premium is aimed at businesses with fewer than 10 employees and costs $12.50 per user per month. The big deal here is simplicity: Microsoft realizes that these businesses have no IT staff, so the company has designed the setup and admin experience in such a way that any employee will be able to get up and running successfully.

Office 365 for Midsized Businesses takes the next step, aiming for 10 to 250 users, and costs $15 per user per month. Building off the small business version, it offers more advanced capabilities around Group Policy management, Windows PowerShell, and federation capabilities, so you can mix and match between the cloud and your on-premises infrastructure. It's still simple, from an admin and deployment standpoint, but more powerful too.

New Focus on Devices and Multi-Touch

Office 365, like Office 2013, has been redesigned to work well with a growing range of devices, most of which will support multi-touch capabilities. This support takes a few different forms, including the full-blown Office applications on Windows, which work well with keyboard and mouse, touch, and even pen input; Office for Mac, which hasn't yet been updated to match the Windows version but supports Apple's trendy desktop/laptop OS; and a growing set of mobile apps that run on Windows Phone, iOS (iPhone, iPad), and Android handsets and tablets.

That last bit is particularly interesting. Windows Phone 8 comes with the latest version of Office Mobile, which includes versions of Word, Excel, PowerPoint, OneNote, and Outlook tailored for the small screens of such devices, and there are SharePoint and Lync apps for the platform as well. Microsoft provides SharePoint and Lync 2013 apps for iOS and will soon do so for Android. And there are even two new Windows 8 "Metro-style" mobile apps, for OneNote and Lync, that perhaps point the way to the Office of the future.

Deeper Social Integration

With this generation of Office products, Microsoft is firmly embracing the social networking wave in several ways. Previously separate Outlook social connectors for services such as Facebook and LinkedIn now ship as part of Outlook 2013, and you can integrate with Skype, Lync, or both—your choice.

On the server and services side, SharePoint's social networking prowess, which was somewhat ignored or misunderstood in some quarters with the previous release, has been expanded and improved. And with Microsoft's Yammer acquisition, the most popular enterprise social networking solution is available via many Office 365 subscriptions at no cost. Deeper integration between SharePoint and Yammer is coming further down the road as well.

Deeper Cloud Integration

As with Windows 8, Office 2013 and Office 365 both straddle the line between work and home, allowing users to utilize familiar interfaces in both contexts. So where the business-oriented versions of Office 365 come with SharePoint Online for document management and collaboration, and integrate seamlessly with Office applications, users are also free to access their own personal documents through SkyDrive. And they can mix and match: In Office 2013, you can configure multiple SharePoint and/or SkyDrive accounts, and move between the different document repositories as needed.

In Windows 7 and Windows 8, users can install local sync solutions for both SkyDrive and SharePoint (the latter sync solution is, unfortunately and confusingly, called SkyDrive Pro) to access those documents offline. And yes, you can install both side by side.

New Features in SharePoint Online

In 2013 guise, the new SharePoint offers a friendlier new web interface, keeping the general Office 365 user experience refresh. SharePoint 2013 supports drag-and-drop file management and other niceties.

The unfortunately named SkyDrive Pro replaces SharePoint MySites as each user's document repository in the cloud. As with SkyDrive, there is a cloud service and PC sync application with the same name, further muddying the waters.

SharePoint is fully extensible, with an apps marketplace whose apps use the same underlying web-based technologies as do Office 2013 apps, so I expect to see a vibrant market for add-ons. SharePoint will be further supported by multiple mobile apps, including one for SkyDrive Pro and one for the social networking aspects of the service, both of which will be available on most platforms.

New Features in Exchange Online

The new Exchange expands on Microsoft's dominant messaging solution with Data Loss Prevention functionality and an Exchange eDiscovery Center for compliance officers. It also adds a gorgeous and useful new version of the Outlook Web App that even supports multi-touch interfaces.

Hybrid deployment options have been added for those organizations that need to keep some data in-house. You'll also find a streamlined Exchange Administration Center.

New Features in Lync Online

Microsoft's latest Office server, Lync, provides a single interface for voice calls, video calls, meetings, presence, and instant messaging, and in 2013 it will expand to include Skype federation functionality, deepening that consumer and business interconnection. (It can already federate with such services as Windows Live Messenger, Yahoo! Messenger, AOL Instant Messenger, and Google Talk.)

Lync already integrates with Outlook and is available in mobile app versions for Windows 8 and Windows RT, Windows Phone 8, iOS, and, soon, Android devices, too. Those who use Lync's video conferencing capabilities will appreciate the change to standards-based H.264 video compression for a higher-quality experience.

A Schedule for Change

With the move to this second generation of Office 365 services, Microsoft is also fine-tuning its plan to deliver updates going forward. Although the company will continue on the quarterly update cycle it established originally for Office 365, there are two major changes to the schedule. First, Microsoft plans to be more aggressive, adding more features, and it is revamping its engineering cycle to accommodate a faster rate of change. Second, this schedule includes, for the first time, the Office 2013 suite of applications as well. So there will be quarterly updates to that software as well, not just bug fixes but new features.

Of course, not everyone will want to move so quickly. Microsoft is allowing customers to roll out this major new update as well as future updates for up to 12 months. And the update is configurable across users, so you can opt to pilot some users on new versions more quickly than others to gather feedback and make sure the changes don't break any existing systems.

Those already on Office 365 were able to move to the new version in late February. This version includes the new admin interfaces, the new web experiences for SharePoint Online, Exchange Online, and Lync Online, as well as the new Office 2013 Professional Plus suite if your subscription supports it.

Overall, Office 365 remains the bellwether for Microsoft's move from traditionally delivered software to online services, and the prognosis so far is overwhelmingly positive. This newly expanded set of subscription services is powerful, affordable, and correctly designed for the markets it addresses. With this latest version, Office 365 has gotten demonstrably better. This is a solution you're going to want to evaluate soon. ■

Get-Credential Gives You Power Over Your Passwords

Save valuable time while running tasks that require different credentials

Mark Minasi

is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 30 books, including *Mastering Windows Server 2008 R2* (Sybex). He writes and speaks around the world about Windows networking.

Email

Twitter

Website

Two months ago, in "3 PowerShell Account Tweaks," I mentioned the PowerShell *get-credential* command in passing, but I didn't really cover it. This month, I want to introduce it to you more fully, because I think you'll find it very useful, particularly if you administer more than one forest or if you often have to do a one-off admin job wherein you don't want to have to log off and back on to get that job done.

Sometimes you're logged on as a domain administrator—but not always. To address that, the Active Directory (AD) PowerShell cmdlets all support a *-credential* parameter. For example, if you're logged on as someone with no domain admin powers, and you run the command

```
set-aduser Tjefferson -title "Prez"
```

that command would fail. Suppose, however, that you do have a domain user account named bigfirm\Kathleen, and you want PowerShell to use that account just for this command. You could then type

```
set-aduser Tjefferson -title "Prez" -credential bigfirm\Kathleen
```

which would cause PowerShell to pop up a GUI logon dialog box with a username text field (pre-populated with *bigfirm\Kathleen*) and a

password text field. Punch in the password, click OK, and the command would run without flaw. If you didn't want to specify the username in the command invocation, you could just type

```
set-aduser Tjefferson -title "Prez" -credential (get-credential)
```

In that case, the same dialog box would pop up—but with an empty username field. You might be wondering how you might put a username and password directly onto the command line, as in

```
set-aduser Tjefferson -title "Prez" -credential
  bigfirm\Kathleen -password domaincrusher
```

Unfortunately, that's not possible. You can, however, save yourself some typing by getting that credential once, saving it, and re-using it. You can do that by running *get-credential* a bit differently.

First, run *get-credential*, but this time, capture its result in a PowerShell variable. Variables are spaces in memory that you can use to store data, and PowerShell identifies variables by their first letter: *$* (the dollar sign). Variables can store just about any kind of data. For example, if you typed the lines

```
$firstnum = 3
$secondnum=10
$myresult= $firstnum + $secondnum
```

into PowerShell and pressed Enter, you'd see a result of 13. Notice that you can give variables any name you want, as long as the first character in that name is *$*. The variables $firstnum, $secondnum, and $myresult didn't exist before I typed them, and when I exit PowerShell, any variables I've worked with disappear. To save myself some typing, I could log on as bigfirm\mark and store that credential into a variable that I'll call *$c* with the command

```
$c = get-credential bigfirm\mark
```

I can then use the credential in a subsequent command, such as the *set-aduser* example, as in

```
set-aduser Tjefferson -title "Prez" -credential $c
```

Now, for the rest of the day, every time I need to run some AD command, I can just add *-credential $c* rather than punching in my password all the time. That's nice, but it gets better if your day requires you to run commands from different accounts, such as running some things under a local account and others under admin accounts from various domains. For example, suppose you need a credential for your local machine on an account named *me* on your computer, *PC429*; one from a domain account bigfirm\alex in one forest; and one from outsiders\wally in another forest. You could put these three commands in your PowerShell profile:

```
$c1 = get-credential PC429\me
$c2 = get-credential bigfirm\alex
$c3 = get-credential outsiders\wally
```

Your system would prompt you for each of those passwords, and you could avoid typing passwords for the rest of the day.

Wondering what's in that credential? Essentially, it's just the name of the account that the credential is based on, as well as a pile of binary data of some type. Ask PowerShell to show you what's in $c by just typing

```
$c
```

PowerShell will respond with the username and inform you that it includes a password, and that the password is a System.Security

.SecureString, which sort of looks like a dead end. A little searching shows that PowerShell has a command called *convertfrom-securestring*. Typing

```
$c.password | convertfrom-securestring
```

will dump a long string of hex, which clearly isn't your password—or is it? When you type

```
$c.getnetworkcredential().password
```

PowerShell will show you the password that you typed in.

In sum, then, you can use *get-credential* to create variables that will store your passwords and save you time when you're running tasks that require different credentials. But be sure to lock your workstation when you walk away from it, or a bit of quick typing might just reveal your password! ■

Free Windows Server Administration Tools

Indispensible tools to help you manage your Windows Server environment

Michael Otey

is senior technical director for *Windows IT Pro* and *SQL Server Pro* and author of *Microsoft SQL Server 2008 High Availability with Clustering & Database Mirroring* (McGraw-Hill).

Email

Microsoft's Windows Server infrastructure has been the IT standard for the past two decades. Not surprisingly, during that time a substantial supporting ecosystem has developed, offering a plethora of tools to help you manage your Windows Server environment. In this column, I'll review 10 of the coolest free tools that are available to help you with your Windows administration. As you might expect, there are more powerful commercial versions of most of these tools. However, for each case, the free version performs a useful function on its own. There are actually so many good free tools that it's difficult to cut the list off at 10. However, here are my top 10 free Windows Server management tools for 2013.

10 Remote Desktop Manager

If you're like me, you do a lot of management using RDP remote desktop connections. It doesn't take long before your desktop is cluttered with all sorts of RDP connections—each one requiring its own settings and a different authentication. Remote Desktop Manager from Devolutions lets you centralize all your remote connections, passwords, and credentials.

9 Enable Remote Desktop Remotely

Windows Remote Desktop is an invaluable remote management troubleshooting tool. However, you need to have remote desktop management enabled on the target computer before you can connect

to it, and not all users are able to accomplish that task easily. IntelliAdmin's Enable Remote Desktop Remotely tool lets you enable RDP remotely.

8 EasyBCD

Working with the old boot.ini file was a piece of cake. However, that changed when Microsoft moved to the Boot Configuration Data (BCD) boot environment in Windows Vista and Windows Server 2008. BCD made the boot process more secure but also more difficult to manage because of its arcane command-line BCDedit tool. EasyBCD from NeoSmart Technologies provides a simple-to-use graphical editor for your Windows BCD boot files.

The Sysinternals Suite is a must-have for Windows administrators.

7 Win IP Config

Like EasyBCD, Win IP Config might be more of a system utility, but if you get tired of running the text-based ipconfig command, you might want to see a graphical equivalent. That's what Win IP Config is. Win IP Config provides the same type of information as ipconfig.exe and route.exe and can also renew IP addresses.

6 Wake-On-LAN

Wake-On-LAN from SolarWinds does just what its name implies. If your networked PCs have Wake-On-LAN enabled in their BIOS, you can send them a packet over the network, causing them to boot up exactly as if you pushed the power button. Wake-On-LAN requires you to input the MAC and TCP/IP addresses of the remote system you want to boot up.

5 Viewfinity Local Admin Discovery

Managing administrator rights can be a problem, and it's possible that sometimes admin rights can be granted and then neglected to be removed. Viewfinity Local Admin Discovery lets you find all users that have been given membership in the local Administrators group.

4 Wireshark

Formerly known as Ethereal, Wireshark is an open-source network protocol analyzer. Sometimes you can encounter tough network questions, and an analyzer such as Wireshark that lets you dig down into the raw network traffic can be a valuable tool to help you answer them.

3 Desktop Central 8

ManageEngine's Desktop Central 8 provides a host of desktop management features, including software deployment, patch management, mobile device management, remote desktop control, service pack installation, and USB device management. ManageEngine offers a free edition for small businesses that can be used to manage up to 25 desktops and 2 mobile devices.

2 Spiceworks

Since it was first released back in 2006, Spiceworks has been adopted by thousands of IT professionals. Spiceworks is oriented toward the SMB and it's the Swiss Army knife of management tools. Spiceworks provides inventory management, monitoring, change management, virtualization management, and IT Help desk support. You can check out our review of Spiceworks 4.5.

1 Sysinternals Suite

Perennially number one on my list of free tools is Sysinternals Suite. If you're new to the Windows platform or to IT, you might not have been exposed to this super-valuable set of tools. The Sysinternals Suite offers an impressive array of tools for tasks such as viewing active TCP network connections, managing open files, viewing opened registry keys, and working with active processes. Many of the Sysinternals tools can be run directly from the website without any installation. The Sysinternals Suite is a must-have for Windows administrators. ■

7 More Reasons to Upgrade to Windows Server 2012 Active Directory

Here's how several key Server 2012 features require, or benefit from, an AD upgrade

Sean Deuby

is technical director for *Windows IT Pro* and *SQL Server Pro* and former technical lead of Intel's core directory services team. He's been a directory services MVP since 2004.

Email

Twitter

Much has been written in these pages about Active Directory (AD) improvements in Windows Server 2012—a lot of it by me! (Check out the Learning Path for that coverage.) IT pros who are responsible for AD need to remember, however, that it's not only explicit AD enhancements, such as the ability to clone virtual DCs, that might drive an AD upgrade.

It might be hard to believe, but AD isn't the center of the IT universe. A more realistic analogy is that it's a critical part of the IT building's foundation, like electricity. Employees must have it, typically 24 × 7, to get anything done, and they don't notice it until it doesn't work. To stretch that analogy a little further, AD also allows you to use a single switch to turn on all the ceiling light fixtures across your office floor instead of one fixture at a time. (This is my *Authenticate once, access many times* analogy.) How does this apply to AD upgrades? A number of new or improved capabilities in Server 2012 require upgrades to AD.

It's been my experience in corporate IT environments that other teams, looking at implementing a new OS or application capability, often aren't aware of the AD upgrade requirements of their project. Or if they are, they don't communicate those requirements to the AD team. When the requirements are finally dropped on the AD team's doorstep, the project plan for the new capability is usually

well underway and the AD folks are forced to devise and implement a hastily planned upgrade project. This phenomenon is a formula for unplanned downtime. The AD team's other alternative is to tell the other team, "*NO*—not for a while." This is not a good career strategy.

Instead, look at the positive aspects of such a situation: If you don't have enough ammunition for a Server 2012 AD upgrade based on direct AD improvements alone, the AD upgrade requirements of other, higher-business-impact features might be enough. However you choose to look at it, it's in your best interest to understand which other Server 2012 features have AD upgrade dependencies. So, which new Server 2012 features require, or benefit from, an AD upgrade?

Dynamic Access Control

Dynamic Access Control (DAC) isn't so much a single feature as it is a combination of new or improved technologies in Server 2012. Together, they provide a higher level of access control to and governance of Windows file servers.

The AD-related requirements for DAC depend on how broadly you plan to implement it:

- If you want to use modern conditional expressions (the ability to AND authorization conditions together with expressions rather than security groups) or Access Denied help, you just need to upgrade your Windows file servers to Server 2012.
- If you want to implement claims in AD and Kerberos, or Central Access Policies (which require claims), you must have one or more Server 2012 DCs (and thus upgrade the schema). You must also configure \Policies\Computer Configuration\Administrative Templates\System\KDC\Domain Controller support for DAC and Kerberos armoring to enable claims support in the Default Domain Controllers GPO.
- If you want to choose either of the two most restrictive policy settings—*Always provide claims* or *Fail unarmored authentication requests*—you must upgrade all the DCs in your domain to

Server 2012 and increase the domain functional level. (See the "Functional Levels" section later in this article.)

Active Directory Federation Services 2.1

Active Directory Federation Services (AD FS) is now a full server role rather than an externally installed capability. If you're using AD FS for federated trust management (and you should be using some kind of federation solution, whether it's AD FS, third-party on-premises software, or IDaaS), version 2.1 has greater flexibility in consuming claims. Before Server 2012, claims were created and stored in AD FS only. These claims were created from the user and group SIDs in the user's Kerberos ticket or from LDAP queries that AD FS made to AD. In Server 2012, AD FS can consume AD user and device claims that are included in Kerberos tickets as a result of domain authentication. This is a much more integrated and flexible situation than that of its predecessor. AD FS includes these claims in a SAML token that it issues for the client, to be used by a web service such as Salesforce.com.

Enabling Kerberos claims in AD FS requires the following:

- DAC enabled and configured with more than one Server 2012 DC, so claims are present in the Kerberos token
- Compound ID (i.e., user + device claims) switched on for the AD FS service account
- Server 2012 AD FS

Active Directory–Based Activation

Product activation has also been integrated into Server 2012 with Active Directory–Based Activation (ADBA). Simply put, ADBA allows a product that supports Generic Volume License Key (GVLK) to automatically activate when it joins a forest that has been activated via the Volume Activation Management Tool (VAMT). This nice step forward for integration is tempered by the fact that the only products that recognize GVLK are Windows 8 and Server 2012. So, your volume activation infrastructure will get simpler . . . eventually. ADBA

requires a Server 2012 schema update (to be able to store activation objects in AD), but it doesn't require any DCs.

DirectAccess

If you aren't already familiar with it, DirectAccess is a Microsoft network technology that enables a domain-joined client with Internet access, anywhere, to appear as though it is actually on the corporate network. Among many improvements to the technology in Server 2012 (and there are many improvements that make DirectAccess easier to deploy and use), you can now domain-join a client when it's off the corporate network. This means you can join and manage a client without it ever having natively touched that network. DirectAccess requires a schema update.

Group Managed Service Accounts

Managed Service Accounts (MSAs)—introduced in Windows Server 2008 R2—attempted to solve the problem of changing service account passwords without impacting the services that relied on them. Unfortunately, a limitation of MSAs is that they must be used on a per-computer, per-service basis. This strongly limited their adoption. Server 2012 introduces Group Managed Service Accounts (gMSAs—have you noticed how lowercase letters are sneaking into acronyms?), which allow services running on multiple hosts to use the same gMSA account. Microsoft SQL Server clusters are a good scenario for gMSAs.

Cross-Domain Kerberos Constrained Delegation

Kerberos Constrained Delegation (KCD) is the ability of a service account to act on behalf of users in multi-tier applications to access a limited (hence "constrained") set of back-end services. Before Server 2012, KCD worked only when the front-end service account was in the same domain as the back-end service. With Server 2012, these multi-tier applications can now span multiple domains.

This capability requires the following:

- At least one Server 2012 DC must reside in both front-end and back-end domains (and thus schema updates in each). The client's domain doesn't need to be upgraded.
- The front-end server hosting the service account must be running Server 2012.

Functional Levels

What about AD domain and forest functional levels? AD-related capabilities in Server 2012 have become much less dependent on these settings than in previous OS versions. The "Functional level features and requirements" section in the TechNet article "Upgrade Domain Controllers to Windows Server 2012" provides more information about current requirements and about functional levels in general.

As I stated earlier, there's one circumstance that might require you to bump the domain functional level. If you want all your DCs to advertise that they support claims (with or without mandatory Kerberos armoring and Flexible Authentication Secure Tunnel—FAST—as defined by RFC 6113), you've got to upgrade them all, then increment the domain functional level.

There are no specific features or capabilities that require Server 2012 domain functional level; it will, however, ensure that if you create any new domains, they will automatically be at the Server 2012 domain functional level.

Start Simplifying

You might not be able to justify Server 2012 AD on its merits alone (though you certainly should implement it), but other improvements might do the trick. Your security team will love gMSAs, and line-of-business (LOB) apps will appreciate the extra flexibility that KCD provides. Also, the basic implementation of AD claims and expressions for access control will let everyone start simplifying their security environment. ■

Learning Path

Learn more about AD improvements in Windows Server 2012:

"Windows Server 2012 Active Directory Moves Forward"

"How Windows Server 2012 Improves Active Directory Disaster Recovery"

"Virtualization-Safe Active Directory in Windows Server 2012"

"Windows Server 2012 Simplifies Active Directory Upgrades and Deployments"

"7 Steps to Clone Windows Server 2012 Virtual Domain Controllers"

How Windows Server 2012 Eases the Pain of Kerberos Constrained Delegation, Part 2

Mike Stephens

is a senior support escalation engineer with Microsoft Commercial Technical Support. He is a Microsoft Certified IT Professional (MCITP-EA) and a Microsoft Certified Trainer (MCT). He recently authored and delivered Windows Server 2012 Distributed Systems Security training for Microsoft Commercial Technical Support. He is a contributor to the Ask the Directory Services Team blog.

Email

In "How Windows Server 2012 Eases the Pain of Kerberos Constrained Delegation, Part 1," I provided an introduction to resource-based Kerberos constrained delegation, included in Windows Server 2012. I also described the targeted scenarios for which resource-based Kerberos constrained delegation is designed, and I provided a brief overview of the technology. Now, in part 2, I want to expand on how resource-based Kerberos constrained delegation works by providing more technical depth as well as a message flow walkthrough.

Diving Into the Technical Depths

Constrained delegation lets you limit the back-end services for which a front-end service can request tickets on behalf of another user. To understand this behavior, it's best to analyze authentication flow as two separate events: the client authenticating to the front-end service, and the front-end service authenticating to the back-end service.

Client to front-end authentication. Authentication from the Kerberos client to the front-end server doesn't change when you use resource-based constrained delegation. The Kerberos client requests a service ticket from its local Key Distribution Center (KDC) for the target service principal name (SPN).

If the target service resides in the same domain, the KDC issues a service ticket and session key to the Kerberos client in a TGS-REP message. If the target service resides outside the current domain, the KDC issues a Ticket Granting Ticket (TGT) referral ticket using the inter-realm session key of the trust in a TGS-REP. The Kerberos

client chases the referral as it normally does when authenticating to a resource outside of its domain (across a trust).

Front-end to back-end authentication. Authentication from the front end (Service-for-User—S4U—client) to the back end is different when using resource-based constrained delegation. Resource-based constrained delegation requires that the computer running the front-end service use Server 2012 because services running on versions of Windows earlier than Windows 8 and Server 2012 don't support resource-based constrained delegation; earlier versions of Windows don't chase referrals from the Service-for-User-to-Proxy (S4U2Proxy) TGS-REQ across the domain boundary.

During front-end to back-end authentication, the front-end service asks a KDC for a service ticket on behalf of another user. This exchange uses the Kerberos extension S4U2Proxy (aka constrained delegation). The Kerberos client successfully presents a service ticket to the front-end service. The front-end service impersonates the identity presented in the service ticket and attempts to authenticate to the back-end service by way of SPN. This authentication attempt results in the front-end service creating an S4U2Proxy TGS-REQ to the KDC in the front-end server's domain. This request includes the target SPN, which resides in another domain, and the service ticket used to authenticate to the front-end service. The TGS-REP returned depends on the answering KDC.

Front-End KDC Behavior

Constrained delegation, at the micro level, involves many decisions and exchanges of information, beginning with the client contacting the front-end KDC.

KDC earlier than Server 2012. A KDC earlier than Server 2012 receiving an S4U2Proxy TGS-REQ for a target SPN outside of its domain returns the Kerberos error KDC_ERR_BADOPTION (13) in a TGS-REP to the front-end service. This response results from an inability of a KDC earlier than Server 2012 to provide a TGT referral

for an S4U2Proxy TGT_REQ for a target SPN residing outside its own domain. Constrained delegation prior to Server 2012 wasn't supported across domain and forest trusts.

Server 2012 KDC. A Server 2012 KDC receiving the S4U2Proxy TGS-REQ determines whether the target SPN resides in its domain. In this example scenario, the target SPN resides in another domain. Therefore, the Server 2012 KDC—aware that it supports resource-based constrained delegation—provides a referral TGT to the front-end service in a TGS-REP.

Front-End Service TGS-REP Behavior

The front-end service receives a TGS-REP from the KDC. The next action the front-end service performs depends on the KDC response from the S4U2Proxy TGS-REP.

TGS-REP from KDC earlier than Server 2012. The front-end service receives a TGS-REP in response to the S4U2Proxy TGS-REQ. The response from the KDC is the Kerberos error KDC_ERR_BADOPTION (13). The front-end service runs on a Server 2012 member server. Server 2012 is a cross-domain constrained delegation–aware Kerberos client; therefore, when the front-end service receives an S4U2Proxy TGS-REP with KDC_ERR_BADOPTION (13), it knows that it might have contacted a KDC that doesn't support constrained delegation across domains. In response, the Server 2012 member server running the front-end service attempts to locate a Server 2012 domain controller (DC). After locating a Server 2012 DC, the member server running the front-end service sends the same S4U2Proxy TGS-REQ to the Server 2012 DC.

TGS-REP from Server 2012 KDC. The front-end service receives a TGS-REP in response to the S4U2Proxy TGS-REQ. The response from the KDC is a TGT referral to the domain that's responsible for providing authentication for the target SPN. Server 2012 is a cross-domain constrained delegation–aware Kerberos client. The member server running the front-end service chases the referral to the domain

listed in the TGT referral. (Important: When traversing trusts using resource-based constrained delegation, the computer must authenticate to traverse the trust. Therefore it is expected for the computer to perform a TGS-REQ for a TGT in each domain as well as the first S4U2Proxy TGS-REQ performed by the front-end service.) The TGS-REQ referral process continues until it locates a Server 2012 DC in the domain that hosts the targeted SPN.

Back-End KDC Behavior

The back-end KDC receives an S4U2Proxy TGS-REQ from the front-end service. The TGS-REQ includes an evidentiary ticket, which is the service ticket from the initial authentication to the front-end service as well as the inter-realm referral TGT received from an earlier exchange with a KDC.

The KDC first determines whether the target SPN resides in its domain. If it doesn't, the KDC creates a referral TGS-REP, as previously described. Alternatively, the target SPN might exist in the current domain. In this case, the KDC can provide a service ticket for the targeted service and can respond directly rather than with a referral to another domain. The KDC then reads the msDS-AllowedToActOn BehalfOfOtherIdentity attribute on the security principal registered for the targeted back-end SPN. If the attribute is empty, the Server 2012 DC will use traditional constrained delegation logic (msDS-Allowed ToDelegateTo [A2D2]). If the msDS-AllowedToActOnBehalfOfOther Identity has a value, the KDC impersonates the security principal under which the front-end service runs and performs an access check using the security descriptor stored in the msDS-AllowedToActOn BehalfOfOtherIdentity attribute.

An access check failure causes the KDC to use traditional constrained delegation logic (A2D2) to determine whether constrained delegation is allowed. A successful access check means the back-end service allows the front-end service to request tickets on behalf of other security principals that are used for authentication to the

back-end service. The KDC builds a service ticket for the back-end service using the client name from the evidentiary ticket and returns the service ticket and session key for the front-end service to use to authenticate to the back-end service as the user.

KDC Behavior With and Without Traditional Constrained Delegation

If the back-end server is configured using traditional constrained delegation (msDS-AllowedToDelegateTo—A2D2), which must reside in the same domain, then a Server 2012 KDC or a KDC running an earlier version of Windows can be used for authentication.

Behavior for non–Server 2012 KDCs. KDCs running earlier versions of Windows behave the same with traditional constrained delegation. If A2D2 isn't configured, and the back-end service resides in the current domain, the KDC returns KDC_ERR_BADOPTION with a substatus of STATUS_NOT_FOUND. If A2D2 isn't configured, and the back-end service resides in another domain, the KDC returns KDC_ERR_BADOPTION with a substatus of STATUS_NOT_FOUND.

If A2D2 is configured, and the back-end service is not a value in the attribute, and the back-end service resides in the current domain, the KDC returns KDC_ERR_BADOPTION with a substatus of STATUS_NOT_FOUND. If the back-end service resides in another domain, the KDC returns KRB-ERR-POLICY with a substatus of STATUS_CROSSREALM_DELEGATION_FAILURE.

Behavior for Server 2012 KDCs. If A2D2 isn't configured, and the back-end service resides in another domain, the Server 2012 KDC returns a referral TGT. If A2D2 isn't configured, and the back-end service resides in the current domain, and resource-based constrained delegation isn't configured on the principal object, the Server 2012 KDC returns KDC_ERR_BADOPTION with a substatus of STATUS_NOT_FOUND.

If A2D2 is configured, and the back-end SPN isn't a value within the attribute, the back-end service resides in the current domain, and

resource-based constrained delegation isn't configured on the principal object, the Server 2012 KDC returns KDC_ERR_BADOPTION with a substatus of STATUS_NOT_FOUND. If the back-end SPN resides in another domain, the Server 2012 KDC returns a referral TGT.

Message Flow Walkthrough

Now that all the academic explanation is out of the way, here's a walkthrough of the message flow to help you visualize how all of this works together. Don't worry if you don't understand all of it the first time! It's a lot to take in, and the changes are a shift in thinking from how delegation used to work to how it can work in Server 2012. The management of it is simple, but the inner workings require a little more thought before they make sense.

To reduce the number of visible steps to those included in the resource-based constrained delegation message exchange, successful client-to-front-end authentication is assumed in Figure 1.

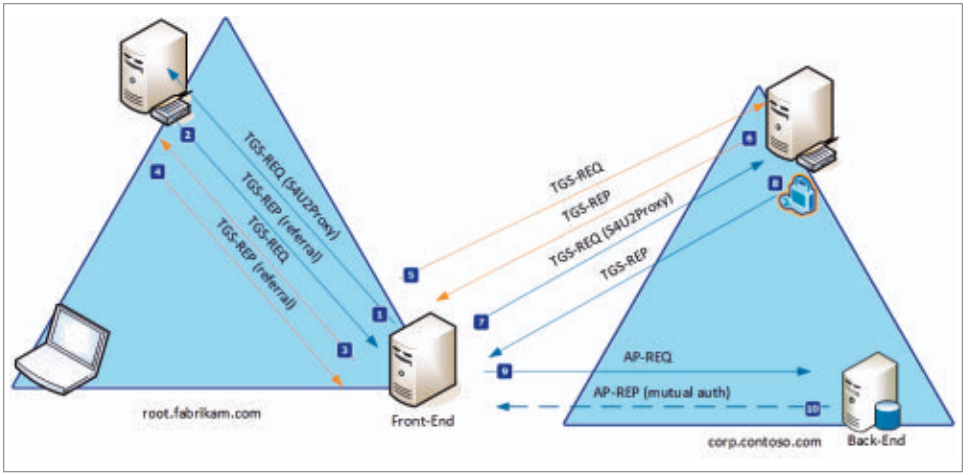

Figure 1 The Resource-Based Constrained Delegation Message Exchange

1. The front-end service sends an S4U2Proxy TGS-REQ to the KDC in root.fabrikam.com, requesting a service ticket for the back-end service on behalf of the user. The TGS-REQ includes the front-end service TGT; a forwardable client service ticket for the

front-end service, or an evidentiary ticket; and the KDC option cname-in-addl-tkt. If the KDC in root.fabrikam.com returns KRB-ERR-BADOPTION, the front-end service locates a Server 2012 DC and retries the TGS-REQ.

2. The KDC in root.fabrikam.com determines that the back-end service doesn't reside in root.fabrikam.com and returns a referral TGT for corp.contoso.com to the front-end service on behalf of the user. The cname field in the ticket uses the name of the front-end service, and the crealm field uses the name of the front-end service domain.
3. The front-end service must authenticate to the back-end domain to chase the referral on behalf of the user. The front-end service sends a TGS-REQ, as itself, to the KDC in root.fabrikam.com to request a service ticket for the back-end service.
4. The KDC in root.fabrikam.com determines that the back-end service isn't in root.fabrikam.com and returns a TGS-REP that includes a referral TGT to corp.contoso.com.
5. The front-end service sends a TGS-REQ, as itself, requesting a service ticket for the back-end service.
6. The KDC in corp.contoso.com sends a TGS-REP that includes a service ticket for the back-end service that is used by the front-end service.
7. The front-end service locates a Server 2012 DC in corp.contoso.com and sends an S4U2Proxy TGS-REQ to the KDC in corp.contoso.com, requesting a service ticket for the back-end service on behalf of the user present in the evidentiary ticket. The request includes a front-end service referral TGT, additional tickets (S4U referral TGT) and the KDC option cname-in-addl-tkt.
8. The KDC in corp.contoso.com retrieves account information from AD using SamIGetUserLogonInformation, impersonates the front-end service, and performs an access check using the security descriptor in the msDS-AllowedToActOnBehalfOfOtherIdentity

attribute. If the access check fails, the KDC returns KRB-ERR-BADOPTION; otherwise, the KDC returns a service ticket in a TGS-REP.

9. The front-end service presents the service ticket requested on behalf of the user to the back-end service by sending an AP-REQ.
10. The back-end service returns an AP-REP if mutual authentication is required.

Protocol Transition (S4U2Self)

The protocol transition extension to Kerberos doesn't require a Server 2012 DC. Therefore, Windows 8 and Server 2012 S4U clients don't attempt to locate a Server 2012 DC to service these requests.

Front-end servers need to locate Server 2012 DCs when the initial S4U2Proxy TGS-REQ returns a KRB-ERR-BADOPTION or KRB-ERR-POLICY. To accomplish this, the S4U client uses the public directory service API DsGetDCName, which makes an RPC call to a DC. The specific call includes the DS_DIRECTORY_SERVICE_8_REQUIRED flag, which indicates the API need only return Server 2012 DCs.

You Asked for It!

This wraps up resource-based constrained delegation, and now you can see how it eases the administrative pains of constrained delegation. Server 2012 simplifies its configuration. It removes the appearance of registering duplicate SPNs on multiple front-end computers, returning the point management to the resource owner rather than the owner of the front-end service and the domain administrator. Also, with resource-based constrained delegation, you can now use constrained delegation across trusted domains and trusted forests—a feature that has been a huge ask from customers for a long time. ■

Using Windows Server 2012 Hyper-V Replica

Asynchronous replication expands your disaster-recovery options

John Savill

is a Windows technical specialist, an 11-time MVP, and an MCSE: Server Infrastructure for Windows Server 2012 and Private Cloud. He's a senior contributing editor for *Windows IT Pro* and his latest book is *Microsoft Virtualization Secrets* (Wiley).

Most organizations completely rely on their IT infrastructures to function. To provide IT resiliency, many ensure that they have backups of their systems. When technically possible, companies also implement high-availability solutions, such as clusters running their Microsoft SQL Server instances and file services, a network load-balanced web farm, multiple domain controllers (DCs) replicating to one another, and so on.

For applications that have no native high-availability capabilities, virtualization can provide a solution by applying high availability at the virtual machine (VM) level. This approach allows the VM to be restarted automatically on an alternative virtualization host if an unplanned failover occurs. It also allows the migration of VMs between hosts, with no downtime, in planned situations such as maintenance events. These solutions handle a failover at the host level (i.e., when one host fails).

However, natural disasters (such as recent "hole in the earth" disasters, referring to the complete loss of the data center) and man-made events (which can be as innocent as road work cutting through both your redundant connections to the Internet) can effectively seal off

your data center from the rest of the world. Organizations must plan for continuing business even if their primary data center is lost.

Assuming that your organization has a second location that can be used as a data center, many of the application-level technologies that I mentioned can also be used across locations. Many solutions, such as multiple DCs and failover cluster–enabled applications, are geographically aware. But there can be a catch with using the traditional Windows Failover Clustering feature over geographically separate locations. In many applications that use clustering and in all virtualization clustering, shared storage must be available to all the nodes in the cluster. This is generally very expensive, because it requires SANs at both locations, great connectivity between the locations, and storage replication to keep the content on both SANs synchronized.

Most SAN-to-SAN replication solutions are synchronous: A write action on the primary SAN is acknowledged to the writing process only when the write is also performed on the replica SAN, thus ensuring that both SANs are synchronized at all times. Although synchronous replication gives the greatest protection, it's costly.

Very large organizations can afford these storage solutions and enable virtualization clusters across locations for top-tier applications and virtual environments that host critical services. But many other organizations and non–top tier applications were left without a way to provide disaster recovery—until the introduction of Windows Server 2012 Hyper-V Replica.

The goal of Hyper-V Replica is to enable disaster recovery capabilities for any Hyper-V environment.

Introducing Hyper-V Replica

Windows Server 2012 was an enormous release, with particularly significant changes around virtualization and cloud services. One of the biggest new features is Hyper-V Replica, which introduces the ability to asynchronously replicate a VM to a second Hyper-V host. The target Hyper-V server (i.e., the replica) doesn't need to be part of a cluster with the primary Hyper-V host. In fact, the replica *can't* be in the same cluster as the primary. Nor does the replica need any shared

storage or even a dedicated network infrastructure for the replication. The goal of Hyper-V Replica is to enable disaster recovery capabilities for any Hyper-V environment, without steep requirements, through its use of asynchronous replication.

Some SANs also offer asynchronous replication, which works by replicating data from the primary to the replica, but not in real time. Write actions are performed on the primary host, acknowledged to the writing process, and then replicated, when possible. There's a delay between when the write occurs on the primary and when it occurs on the replica. Depending on this delay, a certain amount of data can be missing from the replica server—and therefore possibly lost—if the primary host fails. This possible gap is often referred to as the recovery point objective (RPO) and basically defines the maximum amount of data loss that's acceptable in a disaster. For example, an RPO of 5 minutes means that no more than 5 minutes of data should be lost.

SAN-level asynchronous replication might not be desirable to many organizations because it requires the same vendor in both the primary and replica locations. But Hyper-V Replica uses asynchronous replication very efficiently. At a high level, Hyper-V Replica works as follows:

1. When a VM is enabled for replication, a new VM is created on the Hyper-V replica host. This replica VM matches the configuration of the primary VM and is turned off.
2. The storage of the primary VM is replicated to the replica VM on the replica Hyper-V server. A log is started on the primary Hyper-V host to store writes to the replicated virtual hard disks (VHDs). This log file is stored in the same location as the source VHD.
3. After the initial replication of the storage is complete, the log file is closed. A new log file is started to track ongoing changes; the closed log file is sent to the replica Hyper-V host and is merged with the VHDs for the replica VM. The replica VM remains turned off.

4. Every 5 minutes, the log file is closed, a new one is created, and the closed file is merged with the replica.

Hyper-V Replica's use of asynchronous replication opens up the use of replication to many more companies and many more disaster recovery scenarios:

- Data center–to–data center replication for Tier 1 applications in organizations without SAN-level replication, such as small-to-midsized organizations
- Data center–to–data center replication for Tier 2 applications in organizations that have SAN-level replication but don't want to use it for non–Tier 1 applications
- Branch office–to–head office replication, to protect applications that are hosted at a branch location
- Hoster location–to–hoster location replication, for hosting companies
- Replication to a hoster, for disaster recovery at small organizations that don't have a second data center

There are many more potential scenarios. The key point is that with Hyper-V Replica, the ability to replicate VMs is now an option for any organization.

Using Hyper-V Replica

Hyper-V Replica is simple to configure. The easiest way to really understand how Hyper-V Replica works is to walk through its setup options and enable replication for a VM.

The first step is to configure the replica Hyper-V server to accept requests to host a replica. In Hyper-V Manager, choose Hyper-V Settings from the server's list of actions. Within Hyper-V Settings, choose the Replication Configuration list of configurations, as shown in Figure 1. Check the *Enable this computer as a Replica server* check box. You'll then need to make several choices.

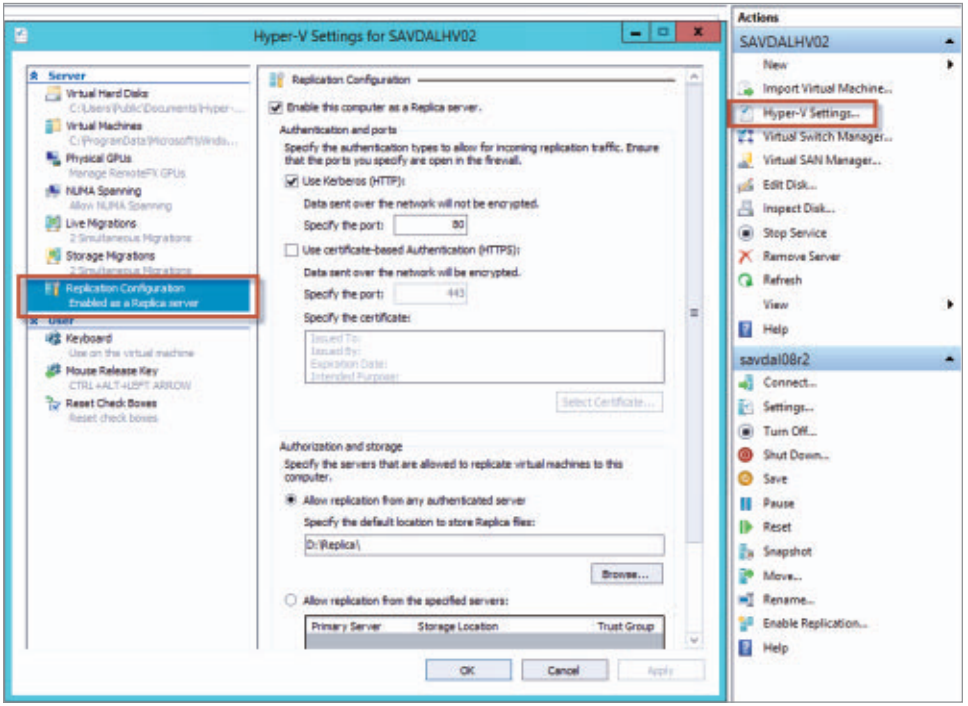

Figure 1
Enabling the Target Hyper-V Server to Accept Replicas

The first choice is to enable the use of Kerberos (which uses HTTP) or certificate-based authentication (which uses HTTP Secure—HTTPS). Kerberos is easier to configure but requires both the primary and replica Hyper-V servers to use Kerberos authentication and therefore be part of the same Active Directory (AD) forest or trusted domains. Using Kerberos, data replication between primary and replica servers isn't encrypted and is sent over the standard HTTP port 80. However, if encryption is required, the Windows IPsec implementation can be used.

The second option is to use certificate-based authentication, which allows the primary and replica servers to be part of different AD forests or organizations. This choice requires a certificate to be specified for use. As an added benefit of using HTTPS, all transferred data is encrypted. If both Kerberos and certificate-based authentication are enabled, then when a new replication relationship is established, the administrator who configures the replication can choose which method to use.

The only other configuration choice is to specify the servers from which the replica will accept replication requests, as well as where those replicas will be stored. One option is to allow replication from any authenticated server. In this case, choose one location to store all replicas. The other option is to specify the servers that can replicate to the replica; each server can have a different storage location.

When specifying servers, you can use one (but only one) wildcard character within the server name. This allows the enablement of a group of servers; for example, *.na.savilltech.net for all servers with a Fully Qualified Domain Name (FQDN) that ends in na.savilltech.net. The Trust Group tag allows VMs to move between Hyper-V hosts with the same trust group and to continue replicating without issue. With Shared Nothing Live Migration, VMs can be moved between unclustered Hyper-V hosts, with no downtime. With this new mobility capability, you need to ensure that groups of servers have the same Trust Group tag to enable unaffected replication when VMs are moved between servers within a trust group.

If you use Failover Clustering, there's an additional requirement. A failover cluster consists of multiple Hyper-V hosts. Therefore, if a failover cluster is the target for Hyper-V Replica, it's important that the whole cluster—not just one host—can host the replicated VM. Therefore, the storage of the replica must be on a Server Message Block (SMB) share or cluster shared volume (CSV). Hyper-V Replica support in a failover cluster is enabled by adding the Hyper-V Replica Broker role to the failover cluster. This action requires a name and IP address for the broker, which serves as the client access point for Hyper-V Replica and will be the name that's used when choosing the cluster as a replication target. When enabling replication within a cluster, you perform the replication configuration within the Failover Cluster Manager tool, after the Hyper-V Replica Broker role is added. When the configurations for replication (which are the same as for a standalone Hyper-V host) are completed, all hosts

in the cluster are automatically configured, unless certificate-based authentication was selected. In that case, each host needs its own configured certificate.

The final step is to enable the required firewall exception for the used port: 80 for HTTP and 443 for HTTPS. The firewall exceptions are built into Windows Server but aren't enabled, even after replication configuration is complete. You'll need to start the Windows Firewall with Advanced Security administrative tool, choose Inbound Rules, and enable either (or both) Hyper-V Replica HTTP Listener (TCP-In) or Hyper-V Replica HTTPS Listener (TCP-In), depending on your authentication method.

When the replica server has been enabled for replication, it's important to also enable the primary Hyper-V server as a replica. This allows the reversal of replication if the VM is activated on the replica server and needs to start replicating back to the previous primary server (which would then be considered the replica).

One item that isn't configured is which network to use for the replication traffic. The assumption is that this technology is used between data centers. There would be only one valid path between them, so Hyper-V Replica automatically chooses the correct network to use for the replication traffic. (I suspect that a number of clients would like more granularity of the network used for Hyper-V Replica; if you're one of them, give Microsoft that feedback!)

Replicating a VM

After the Hyper-V hosts and clusters are configured to enable the Hyper-V Replica capability, the next step is to enable VMs to be replicated. Use Hyper-V Manager or Windows PowerShell (particularly in any kind of automated, bulk configuration). Choose the VM on which you want to enable replication, and then choose the Enable Replication action. This action launches the replication-configuration wizard, which comprises several steps. I walk through the whole process in the accompanying video.

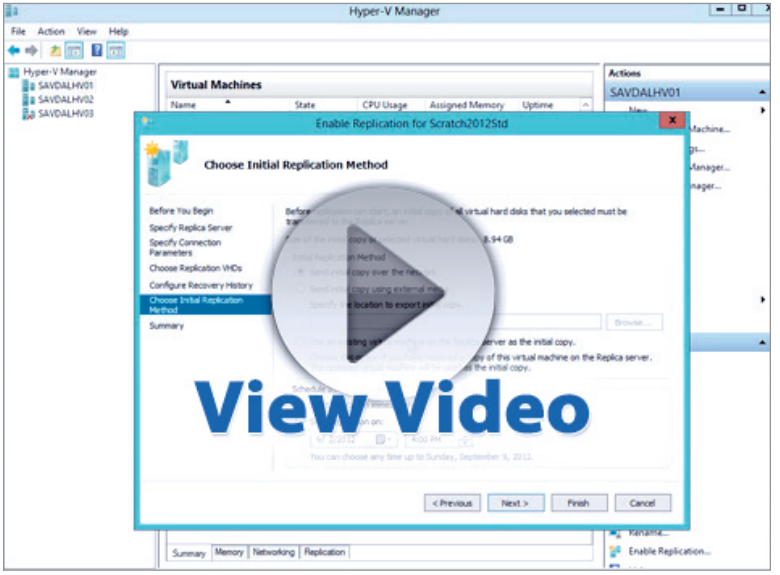

Video

John Savill demonstrates how to configure Windows Server 2012 Hyper-V Replica

After the target Hyper-V server is specified, choose the authentication type to use. This will depend on which types the replica server supports. Also choose whether to compress the data that's sent over the network; compression saves network bandwidth but uses additional CPU cycles on the primary and replica Hyper-V servers. If a VM has multiple VHDs, you can choose which hard disks to replicate. You can use this choice to ensure that only the required VHDs (e.g., only VHDs that contain more than one pagefile) are replicated. Only VHDs can be replicated; if a VM uses pass-through disks, those disks can't be replicated with Hyper-V Replica (another reason to avoid pass-through disks).

The next configuration step is to configure the recovery history. By default, the replica has a single recovery point: the most recent replication state. However, an extended recovery history can be configured to include additional hourly recovery points, as shown in Figure 2. These additional points are manifested as snapshots on the VM that's created on the replica server. You can then choose a specific recovery point by choosing the desired snapshot. An additional option lets you create an incremental Microsoft Volume Shadow Copy Service (VSS)

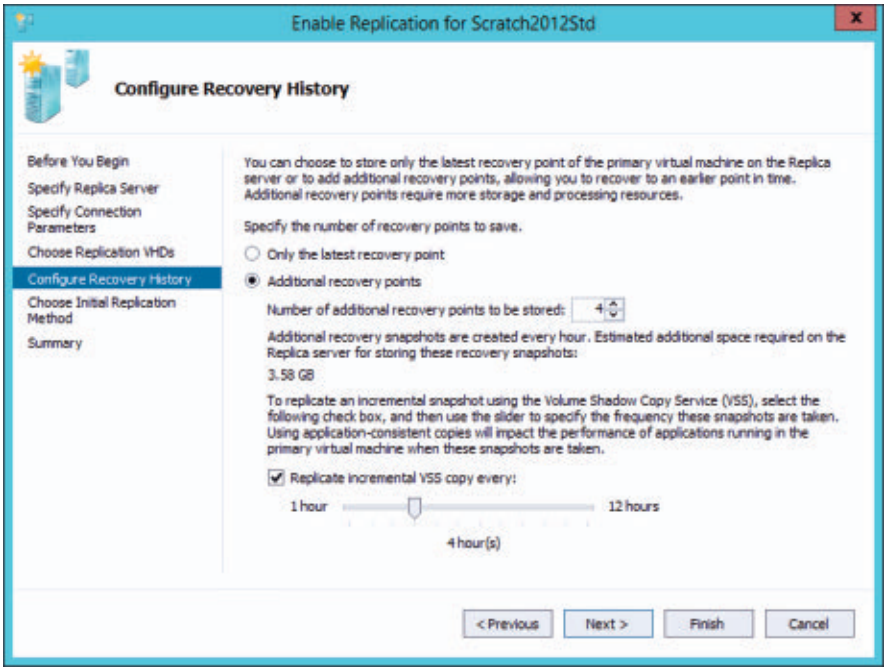

Figure 2
Configuring Recovery History for the VM Replica

copy at a configurable number of hours. This gives you an additional level of assurance in the integrity of the replica at that point in time. The normal log files that are sent every 5 minutes provide the latest storage content. However, at that point, the disk might have been in an inconsistent state on the source VM. There's no guarantee that the replica VHD will be in a consistent state when the replica is started. When enabled, the incremental VSS option triggers a VSS snapshot on the source prior to that cycle's replication, which forces the source VM to ensure that the disk content is in an application-consistent state. In the same manner as when a backup is taken and the log file closed and sent to the replica, that state is saved as the application-consistent recovery point on the target, as shown in Figure 3. If the VM contains applications that have VSS writers, I suggest using the option to create an application-consistent recovery point. The default of 4 hours is a good balance between integrity and the additional work caused by creating a VSS recovery point on the source VM.

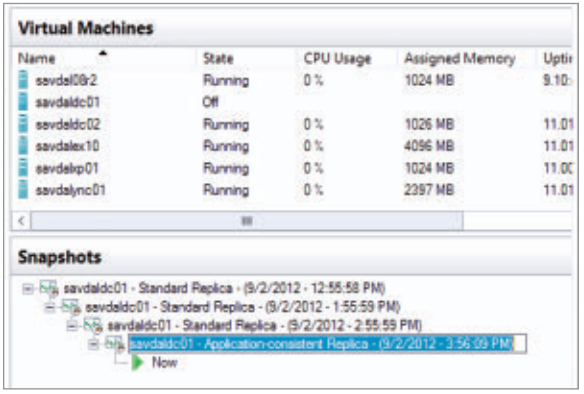

Figure 3
Viewing the Replica

After the recovery-point configuration is complete, choose the method to initially replicate the storage:

- Send VHD content over the network.
- Send VHD content via external media; specify an export location.
- Use an existing VM on the replica server as the initial copy. You can use this option if you already restored the VM to the target Hyper-V server or previously had replication enabled and broke the replica but now want to re-enable it. A very efficient bit-by-bit comparison will be performed between the primary and replica, to ensure consistency.

The initial replication can be configured to begin immediately or at a later, specified time; for example, outside of normal business hours, when contention for network resources is reduced. Depending on your choices, the VM is created on the replica Hyper-V server in the off state, and the initial replication begins. Every 5 minutes, the Hyper-V Replica log (.hrl) file is closed, sent to the replica, and merged into the replica VHD. The entire time, the replica VM is turned off. Only disk content—not memory, processor, or device state—is replicated to the replica VM. If the replica is activated, it will be turned on and booted similar to a crash-consistent state, as if it had just been powered down without clean shutdown. This is one of the reasons why performing the periodic VSS snapshot recovery point is useful for ensuring disk integrity.

After the replica VM is created, it's separate from the primary VM. Any changes in configuration to the primary VM aren't reflected in the replica VM. This allows changes to be made on either VM, and the replication of the VHD content will continue.

Using Hyper-V Replica

Remember that Hyper-V Replica is a disaster-recovery solution. It isn't designed to be used in place of failover clusters or other high-availability technologies. Typically, during a disaster, many steps and processes must be performed to activate a disaster-recovery site. Hyper-V Replica isn't an automatic solution. It won't detect that the primary VM host is missing and start the VM on the replica server because incorrectly detecting a site failure could cause a huge problem. This out-of-the-box feature must be initiated manually, but there's no reason that it can't be automated through PowerShell as part of your other processes. (Perhaps in the future, the feature will be automated through some Microsoft system management solution, such as System Center Virtual Machine Manager, to allow multiple VMs to fail over as part of a larger site-recovery process.)

There are three types of Hyper-V Replica failover—one for testing purposes and two for real disaster scenarios:

- Test failover—This type of failover is triggered on the replica VM. The replica VM can then be started on the replica Hyper-V host. To do so, create a temporary VM that's based on the selected recovery point, and then test to ensure that replication is working as planned and as part of a larger site-failover test process. During the test failover, the primary VM continues to send log updates to the replica VM. These updates are merged into the replica VHDs, ensuring that replication continues. When testing is complete, the temporary VM is deleted
- Planned failover—This type of failover is triggered on the primary VM and is the preferred failover type. This process shuts down the primary VM, replicates any pending changes to ensure that no data is lost, fails over to the replica VM, reverses the replication so that changes flow in the reverse direction, and then starts the replica VM. That VM becomes the primary, whereas the old primary becomes the replica.

- Unplanned failover—This failover type is triggered on the replica VM, the assumption being that in an unplanned failover the primary is unavailable. When this type of failover is performed, a replication of pending changes isn't possible, and reverse replication must be manually enabled with a resynchronization because there's no way to know at which point replication stopped. When starting the reverse replication, choose *Do not copy the initial replication* on the Initial Replication page. The original primary VM can be used, and a block-by-block comparison is performed to synchronize between the replica VM and the original primary VM. Only the delta content needs to be sent over the network.

Something might be bothering you about the failover to the disaster-recovery site in a different location: The VM has a TCP/IP configuration that's unlikely to work in a separate location, which will almost certainly be on a different subnet. As part of the Hyper-V Replica functionality, an additional Failover TCP/IP configuration is available on the VM when replication has been enabled. This configuration allows an alternative IPv4 or IPv6 configuration to be specified on the replica VM. The network configuration is injected into the VM during a failover, as shown in Figure 4.

This process works by Hyper-V updating the VM through the Windows Server 2012 Hyper-V integration services running inside the VM. The process works only on synthetic network adapters, not on legacy network adapters, and requires Windows XP Service Pack 2 (SP2) or Windows Server 2003 SP2 and later to be running inside the VM. At the time of writing, this process doesn't work with Linux VMs but is actively being worked on, so that functionality should be available soon. A good practice is to complete the Failover TCP/IP configuration on the primary VM with its normal IP configuration. That way, if the replica is ever activated, replication reversed, and the VM failed back to the original primary, the correct IP address for the primary location can automatically be reinstated.

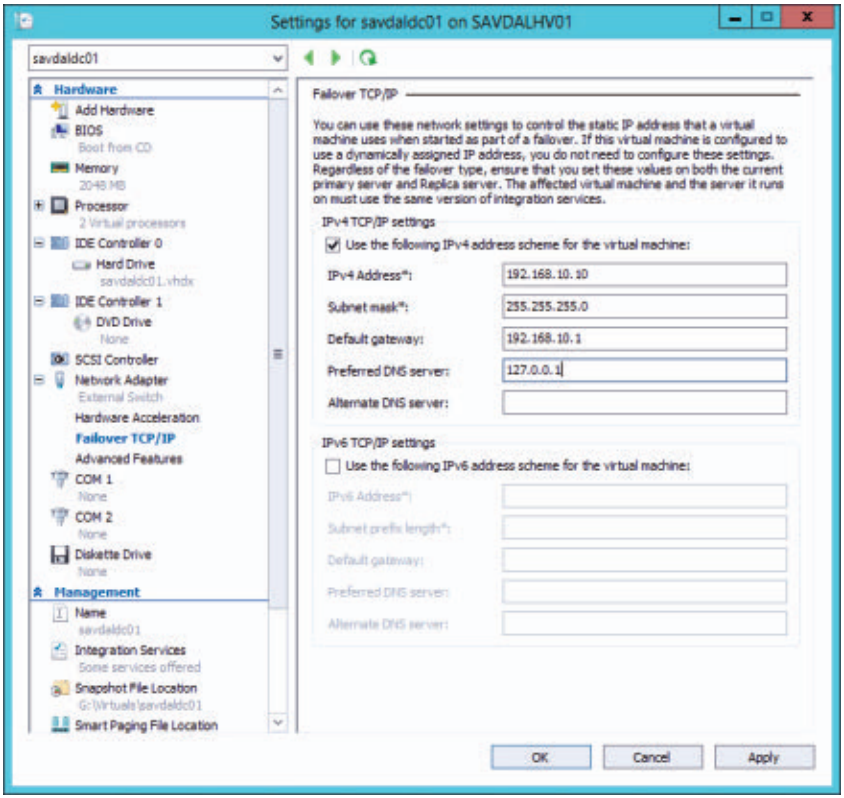

Figure 4
Specifying Alternative IPv4 Configuration for a VM

Replication for Recovery

Hyper-V Replica is a powerful feature. I teased earlier that it's useful even for organizations without a second data center; remember, certificate-based authentication is possible with replication over HTTPS. If you have a hoster that supports Windows Server 2012 Hyper-V (or hopefully Windows Azure Infrastructure as a Service—IaaS), you can replicate from your data center to the public cloud for disaster-recovery purposes. On its own, Hyper-V Replica is a great way to enable failover for individual VMs, but this functionality can also be used by other processes and orchestration components to quickly provide a powerful site-recovery feature that will benefit most organizations. ■

Get Hex Dumps of Files in PowerShell

Use this script to work some magic

A number of years ago, I had an old MS-DOS utility (the exact name of which I can no longer remember) that let me perform a hexadecimal dump of a file. That is, the utility output the hex value of each byte in the file and the ASCII characters for each byte (if printable). This kind of output can be helpful for examining the actual contents of a file, byte by byte. For example, you can see printable strings inside binary files, and you can examine the bytes of a text file.

Some programs can open a file in binary or hex mode. Figure 1 shows an example of a file opened in a text editor that offers hex mode. The gray area on the left is the offset within the file (in hex); the bytes in the file are shown as hex values, 16 bytes per row. The right-hand side of the figure displays the printable ASCII characters for each hex value (unprintable characters are displayed as dots).

I wanted to have this kind of output available in Windows PowerShell. That way, I could quickly perform a hex dump of a file from the command line, without needing a separate program.

```
00000 4D 5A 90 00 03 00 00 00  04 00 00 00 FF FF 00 00  MZ..............
00010 B8 00 00 00 00 00 00 00  40 00 00 00 00 00 00 00  ........@.......
00020 00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00  ................
00030 00 00 00 00 00 00 00 00  00 00 00 00 E8 00 00 00  ................
00040 0E 1F BA 0E 00 B4 09 CD  21 B8 01 4C CD 21 54 68  ........!..L.!Th
00050 69 73 20 70 72 6F 67 72  61 6D 20 63 61 6E 6E 6F  is program canno
00060 74 20 62 65 20 72 75 6E  20 69 6E 20 44 4F 53 20  t be run in DOS
00070 6D 6F 64 65 2E 0D 0D 0A  24 00 00 00 00 00 00 00  mode....$.......
00080 83 C2 32 29 C7 A3 5C 7A  C7 A3 5C 7A C7 A3 5C 7A  ..2)..\z..\z..\z
00090 CE DB D8 7A C6 A3 5C 7A  CE DB C9 7A C5 A3 5C 7A  ...z..\z...z..\z
000A0 CE DB CF 7A DA A3 5C 7A  C7 A3 5D 7A 33 A3 5C 7A  ...z..\z..]z3.\z
000B0 CE DB DF 7A D3 A3 5C 7A  CE DB D5 7A CC A3 5C 7A  ...z..\z...z..\z
000C0 CE DB C8 7A C6 A3 5C 7A  CE DB CD 7A C6 A3 5C 7A  ...z..\z...z..\z
000D0 52 69 63 68 C7 A3 5C 7A  00 00 00 00 00 00 00 00  Rich..\z........
```

Figure 1
Binary File Opened in Hex Mode

Bill Stewart

is a scripting guru who works for Indian Health Service in Albuquerque, New Mexico. He's a contributing editor for *Windows IT Pro* and a moderator for Microsoft's Scripting Guys forum. He offers free tools on his website.

Email

Website

Using Get-Content to Dump a File

It's possible to dump a file, byte by byte, in PowerShell, using the Get-Content cmdlet's -Encoding Byte parameter. For example, the command

```
Get-Content C:\Windows\notepad.exe -Encoding Byte
```

outputs the Windows Notepad.exe program file as an array of bytes. This is a good start: We're getting the bytes from the file. But I want to convert this array of bytes into a hex dump view, similar to the one in Figure 1.

One of the first problems I noticed was that this command can perform slowly, even on relatively small files. This is because Get-Content reads the entire file into memory before it outputs the file as an array.

For improved performance, Get-Content offers the -ReadCount parameter, which lets you specify a buffer size. For example:

```
-ReadCount 16
```

means that Get-Content will read the file 16 bytes at a time (i.e., as a series of 16-byte arrays). If the file's size isn't a multiple of 16 bytes, then the final array will contain fewer than 16 bytes. You access each array by using ForEach-Object and the $_ variable. For example, the PowerShell code in Listing 1 outputs Notepad.exe in hex format, with each byte represented as a pair of hex digits, 16 bytes per line.

Get-Content's -ReadCount parameter certainly improves performance, but I noticed in my testing that the repeated string concatenations (at Callout A in Listing 1) slowed the script as the size of the file

Listing 1: GetHex1.ps1

```
Get-Content "C:\Windows\notepad.exe" -Encoding Byte `
  -ReadCount 16 | ForEach-Object {
  $output = ""
  foreach ( $byte in $_ ) {
A    $output += "{0:X2} " -f $byte
  }
  $output
}
```

increased. I wanted to find out whether I could improve performance by using a different method.

A Faster Method

To read the file, I decided to use the .NET Framework System.IO.File object's OpenRead method instead of Get-Content. The OpenRead method returns a read-only FileStream object. The FileStream object's Read method can read a file a certain number of bytes at a time, similar to the Get-Content cmdlet. But rather than limit the buffer size to 16 bytes, I decided to use a larger buffer and step through it 16 bytes at a time. By using this technique, I was able to avoid string concatenations, except for the last bytes of the file (if the file's size isn't a multiple of 16 bytes). This technique is shown in Listing 2.

The script in Listing 2 opens Notepad.exe as a read-only FileStream object. The script then creates a 64KB byte array, which acts as a buffer for the FileStream object's Read method, as shown at Callout A in Listing 2. The Read method returns the number of bytes it retrieved from the file. Next, the code steps through the buffer in

Listing 2: GetHex2.ps1

```
$bufferSize = 65536
$stream = [System.IO.File]::OpenRead(
  "C:\Windows\notepad.exe")
while ( $stream.Position -lt $stream.Length ) {
A   $buffer = new-object Byte[] $bufferSize
    $bytesRead = $stream.Read($buffer, 0, $bufferSize)
    for ( $line = 0; $line -lt [Math]::Floor($bytesRead /
    16); $line++ ) {
      $slice = $buffer[($line * 16)..(($line * 16) + 15)]
      (("{0:X2} {1:X2} {2:X2} {3:X2} {4:X2} {5:X2} ") +
      ("{6:X2} {7:X2} {8:X2} {9:X2} {10:X2} {11:X2} ") +
      ("{12:X2} {13:X2} {14:X2} {15:X2} ")) -f $slice
    }
```

Listing 2: *continued*

```
B  if ( $bytesRead % 16 -ne 0 ) {
     $slice = $buffer[($line * 16)..($bytesRead - 1)]
     $output = ""
     foreach ( $byte in $slice ) {
       $output += "{0:X2} " -f $byte
     }
     $output
   }
}
$stream.Close()
```

16-byte increments. The code uses the Math object's Floor method to find out how many 16-byte chunks are in the buffer (the final call to the Read method might return less than 64KB). Using the -f operator, the script in Listing 2 takes a 16-byte slice of the buffer and outputs these 16 bytes as hex digits.

If the Read method retrieved a number of bytes from the buffer and that number isn't a multiple of 16, then the *if* expression in the first line at Callout B in Listing 2 returns a non-zero value. In this case, the code in Listing 2 uses string concatenation, similar to the code in Listing 1, to output the file's final bytes.

Both Listing 1 and Listing 2 produce identical output. However, Listing 2 provides a performance improvement over Listing 1 because the repeated string concatenation for the $output variable happens only once for the final bytes at the end of the file, and only if the file's size isn't a multiple of 16.

There are two things missing from these sample scripts: The file offset (the gray area in Figure 1) and the ASCII representation of each byte (the right-most 16 characters in Figure 1). All I needed to do was add these two enhancements, and my script was ready for prime time. The completed script Get-HexDump.ps1, which you can download, includes these features (and some other enhancements as well).

Download

Download the code

Using Get-HexDump.ps1

The script's command-line syntax is as follows:

```
Get-HexDump.ps1 [-Path] <String> [-UnprintableChar <Char>]
  [-BufferSize <UInt32>]
```

The -Path parameter specifies the name of a file. Because this parameter is the script's first positional parameter, the -Path name itself is optional. Wildcards aren't permitted: The script can dump only one file at a time.

The -UnprintableChar parameter specifies the output character to use for characters that aren't in the standard ASCII printable range (characters 32 through 126). The default character is a dot (.). You can use this parameter to specify a different character. If you want to use a space, specify a single space character inside single (' ') or double (" ") quotation marks.

The -BufferSize parameter lets you specify the buffer size that the script uses for reading the contents of the file. The default buffer size is 65,536 bytes (64KB). The -BufferSize parameter's argument must be a multiple of 16.

For each 16 bytes of a file, Get-HexDump.ps1 outputs a string containing the following:

- the offset within the file, in hex
- the hex values of the 16 bytes
- the ASCII character representation for each printable byte
- a placeholder character (a dot or whichever character you specify for the -UnprintableChar parameter) for each unprintable byte

The script does the following "sanity checks" before opening the file:

- makes sure that the file exists
- checks that the file is less than 4GB (the file offset in the output goes only to 0xFFFFFFFF)
- verifies that the requested buffer size is a multiple of 16

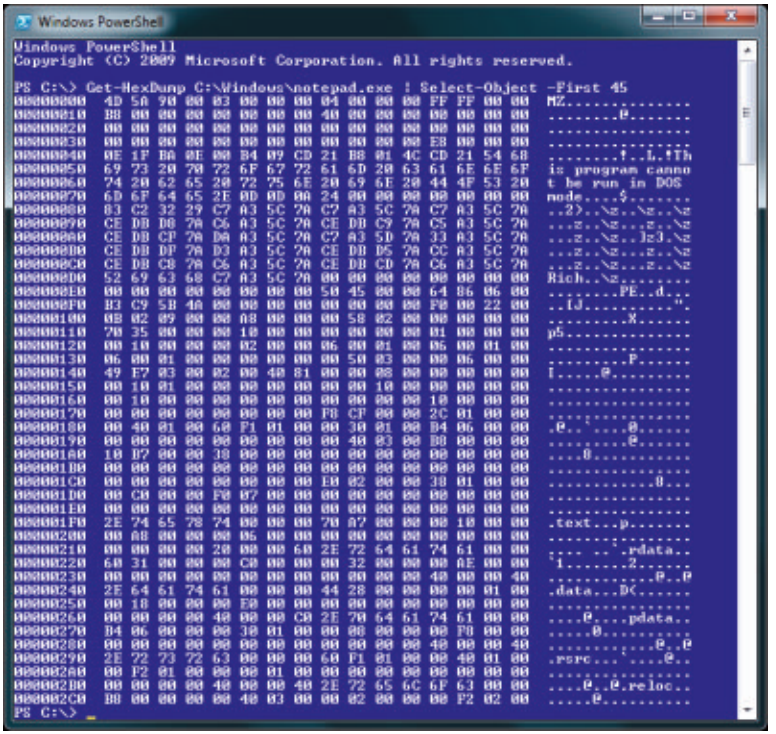

Figure 2
Get-HexDump.ps1 Output of First 720 Bytes of Notepad.exe

The script uses try/catch/finally blocks for error management. The main body of code is enclosed within the *try* block. If a terminating error occurs, the *catch* block outputs the error object. The *finally* block closes the file, regardless of whether an error occurred. The script also uses the Write-Progress cmdlet to visually indicate its progress, which is helpful when redirecting the script's output to another command or file.

You can experiment with the -BufferSize parameter to assess the performance impact on the script. (The Measure-Command cmdlet is helpful here.) Note that even with a much larger buffer, performance won't increase dramatically because PowerShell is still outputting a formatted string every 16 bytes. A small buffer will make the script run more slowly. The script updates its progress bar (using Write-Progress) after each buffer read, so the progress bar will move more slowly, with a greater percentage completed each time, as you increase the buffer size. After some experimentation, 64KB seemed like a reasonable default buffer size. Figure 2 shows an example of the Get-HexDump.ps1 script's output of the first 720 bytes (i.e., 45 × 16) of C:\Windows\Notepad.exe on my Windows 7 x64 SP1 system.

One Less Limitation

PowerShell doesn't have a native cmdlet suitable for viewing the content of binary files. With the Get-HexDump.ps1 script in your toolbox, this is a limitation you no longer need to live with. ■

Managing Exchange ActiveSync Device Access

Allow, block, or quarantine?

Since Microsoft introduced Exchange ActiveSync (EAS) as part of the old Microsoft Mobile Information Server, it has steadily become more widespread. A combination of smart business decisions and technical capability has made EAS the most commonly deployed mobile-device access protocol for Microsoft Exchange Server.

EAS differs from IMAP and POP by combining a protocol for mail access with device management and control. The EAS protocol specifications explain how to retrieve and send mail, perform searches, fetch and apply policies, and so on. (Fun fact: Do an Internet search for "EAS protocol" and you'll find more than you wanted to know about the US Government's Emergency Alert System; but I digress.)

Three EAS-related concepts form the basis of the solution (and understanding how to make it do what you want):

- The protocol specification explains how devices and Exchange servers are supposed to talk to one another: which commands each can emit, which kinds of responses are legal for each command, and so on.
- Exchange implements EAS on the server side. Besides the code that actually sends and receives data using EAS, there's code that allows administrators to view and set EAS policies through the Exchange Management Console (EMC) and Exchange Management Shell (EMS), code for logging, and code for controlling which devices can connect and what they can do after connecting.
- Clients implement EAS, too. Their implementations are supposed to allow the clients to send and receive email, manipulate existing

Paul Robichaux

is a senior contributing editor for *Windows IT Pro*, a content author at Acuitus, and a Microsoft Exchange Server MVP and MCSE. Paul is the author of *Exchange Server Cookbook* (O'Reilly and Associates).

Email

Blog

email items (with operations such as deleting or flagging messages), and accept and enforce policy settings that come from the server.

Exchange Server 2010 Service Pack 2 (SP2) implements EAS version 14.2. However, not every client can say the same. Because so many mobile devices support EAS, you can't assume that any particular device family supports a particular EAS version, or even all the features in that version. Microsoft started a Wiki page to document which specific EAS features various device and OS combinations support. It's a valuable reference because it tells you whether a given device will support an EAS policy or feature that you need. For example, Apple's iOS 5 and Google's Android Ice Cream Sandwich release both support the *disable camera* EAS setting, but other versions don't. These policy controls are useful, but what do you do to keep out devices that don't implement the policies that you want? That's what the Allow/Block/Quarantine (ABQ) functionality is intended to do.

EAS Policies

Before we talk about ABQ, we need to explore a few EAS policy concepts a bit more. When you install Exchange 2010, you get a default EAS policy that doesn't apply anything in the way of security controls. The policy is there just so that when new devices sync with the server, there's a policy to give them.

Many protocols provide some means of negotiation, whereby the client and server can agree on exactly how their communications should go. EAS isn't one of these protocols. When a device first synchronizes, the server and client negotiate which EAS version they'll use; they agree on the highest version that they both support. This happens when the client sends a provisioning request to the server and the server responds with a temporary synchronization key and a policy. The client should return an indication that it will enforce the server-supplied policy. The client then gets a permanent synchronization key, and there's no further checking or enforcement of what the

client says. In an EAS conversation, the Exchange server presents a policy to the client, and the client is expected to honor it.

A mechanism exists to allow the client to indicate that it couldn't apply some parts of the policy in either of two ways: The client can indicate that part of the policy is irrelevant (e.g., disabling Bluetooth on a device that doesn't support it) or that the policy couldn't be applied (e.g., requiring device encryption that isn't supported). Beyond this mechanism, the server has no enforcement mechanism to verify that the client is telling the truth. For example, some versions of Apple iOS lie about whether they can enforce device encryption when the EAS policy is set to require it. Therefore, you might create a policy that's honored by only some of the devices that sync with your server, and you won't know unless you keep track of which devices you have and the specific policy features that they do and don't support.

The simplest way to apply a given set of EAS settings is to change the settings of the default policy to match the values that you want. You can create additional policies to assign to users or groups; each user mailbox can have zero or one EAS policy associated with it. (For more information about managing EAS policies, see "Managing Exchange ActiveSync Policies in Exchange 2010.")

This situation has led to the current ABQ feature set in Exchange 2010. Rather than expect the client to be truthful about its implementation, Microsoft has given us a set of tools to keep out clients whose EAS implementation doesn't conform to a given policy's requirements.

Exemptions, Rules, and Device Access States

One somewhat confusing aspect of Exchange 2010 is that you can sometimes apply a given setting in more than one way. This is definitely true of mobile-device access. You can allow or deny users mobile access to their mailboxes in three ways.

First, you can opt to use the Set-CASMailbox cmdlet with the -ActiveSyncEnabled parameter to turn off EAS access for the user's mailbox. Of course, if you turn off EAS, the user won't be able to

synchronize any device (which might be exactly what you want). As a security measure, many organizations disable EAS for all mailboxes and then enable it only for users who have permission to sync their devices.

Second, you can create explicit personal exemptions that either allow or block devices on a given mailbox. For example, you might set the CEO's mailbox so that she can sync her iPad and iPhone, but nothing else. These exemptions are set by using the Set-CASMailbox cmdlet with the -ActiveSyncAllowedDeviceIDs and -ActiveSyncBlockedDeviceIDs switches to allow and block devices, respectively. The device ID is unique to a given device; think of it as a globally unique identifier (GUID) specific to an individual gadget. The simplest way to get these IDs is to allow the device to synchronize and then use Get-CASMailbox, like this:

```
Get-CASMailbox | select name, ActiveSyncAllowedDeviceIDs
```

That pipeline gives you a table of all the currently enabled devices and their IDs. If you just want the device IDs associated with a single mailbox, you can use the Get-ActiveSyncDeviceStatistics cmdlet, as shown it the following example:

```
Get-ActiveSyncDeviceStatistics -mailbox "paulr"
```

The third way to allow or deny access is to create device access rules. This name is a little misleading; when most of us think of Exchange rules, we think of things such as Outlook rules or transport rules, which typically contain a condition, an action, and a set of exceptions. EAS ABQ rules specify a device family and a device model. The predefined families include the iPhone, iPad, iPod, Android, and Windows Phone. After a user has synchronized a device that you want to use as the basis for a device access rule, you can create one by performing the following steps:

1. Open Exchange Control Panel (ECP) and launch the Manage My Organization page.
2. Choose *Users and Groups*, find the user with the target device, and open the user properties.
3. Expand the Phone & Voice Features item, choose Exchange ActiveSync, and click Edit.
4. Choose the device for which you want to create a device access rule, then click *Create a rule for similar devices*, as Figure 1 shows.

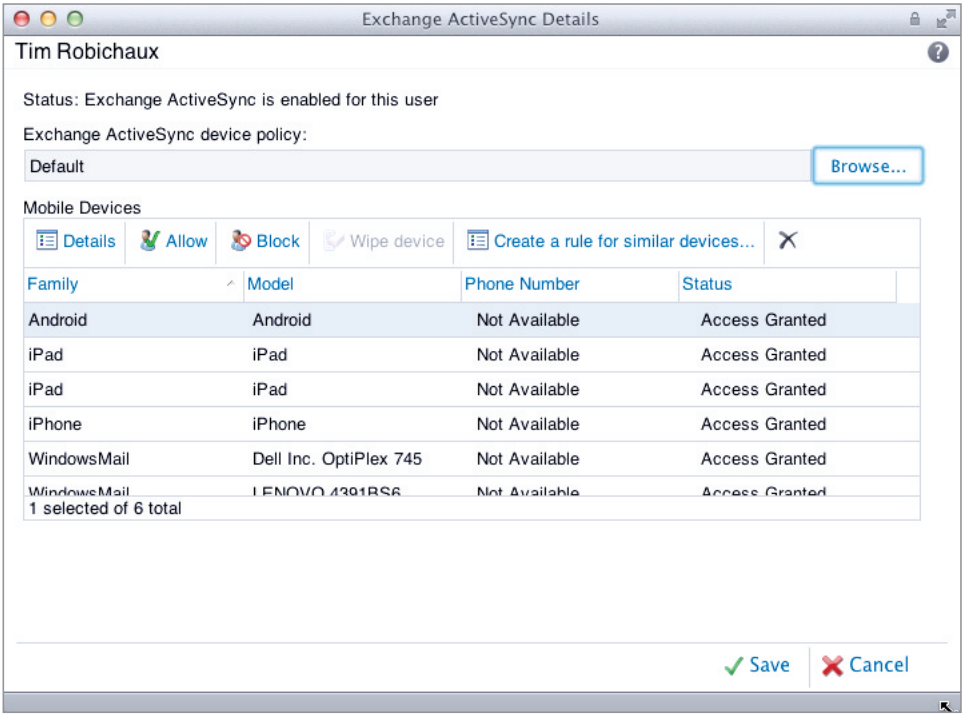

Figure 1
Choosing a Device to Use as a Rule Basis

You can also manipulate device access rules from EMS by using the *-ActiveSyncDeviceAccessRule cmdlet. You'll need to do this if you want to create access rules for devices that don't have family or device strings already created. Check the documentation for New-ActiveSyncDeviceAccessRule and you'll see how it's done: You

specify the device name, the UserAgent string that the device delivers when it makes HTTP requests, the device model, and so on. Be aware, though, that the values you use here must exactly match what the device reports, so you might need to let a test device sync so that you can see which model, UserAgent, and so on the device actually reports.

The ABQ mechanism also relies on the fact that every EAS device that's associated with a server is in one of five possible access states:

- In the *device discovery* state, the device has connected and requested synchronization. The device doesn't actually synchronize anything; the server treats the device as though it were quarantined.
- In the *allow access* state, the device is permitted to access and synchronize with the mailbox.
- In the *block access* state, the device receives an HTTP 403 Forbidden error. The user might receive an email message indicating that the device is blocked if the block is caused by an explicit ABQ setting. If the device is blocked because it didn't apply policies correctly or because it's a non-provisionable device, then the user isn't notified.
- In the *quarantine access* state, the device can't read any mailbox data, but it can post new calendar appointments, contacts, tasks, and notes to the mailbox. After the first sync attempt from a quarantined device, the user receives one email message indicating that the device is quarantined. This message appears both on the device and in the user's Inbox. You can configure a set of administrators who should receive email when a new device enters this state.
- The *mailbox upgrade* state is used when you move a mailbox from an older version of Exchange to Exchange 2010. A device that's associated with a moved mailbox can sync for as many as 7 days. If the device doesn't enter the allow, block, or quarantine state during that time, it loses all access.

Microsoft has given us a set of tools to keep out clients whose EAS implementation doesn't conform to a given policy's requirements.

How do devices move between states? All new devices (i.e., those that have never previously tried to sync with a mailbox in

the organization) start in the device discovery state. This state is short-lived, though. The longest that a device can stay in the device discovery state is 14 minutes (1 minute less than the default EAS heartbeat timeout interval). After that, EAS uses a straightforward algorithm to determine which state to use. This algorithm technically begins with forcing the device to authenticate. (If the device can't authenticate, then it doesn't enter any state because the Client Access server won't talk to it.) After the device authenticates, here's how the algorithm works:

1. If EAS isn't enabled for the user's mailbox, the Client Access server returns an error to the device, and synchronization fails.
2. If the device doesn't meet the criteria for applying and enforcing the policy, it's blocked, and synchronization fails.
3. If the device is blocked by an explicit personal exemption, synchronization fails.
4. If the device is explicitly allowed by a personal exemption, it's granted full access.
5. If the device is blocked or quarantined by a device access rule, its state is changed accordingly and synchronization stops. (Note that the term "fails" would be inaccurate because no error is returned to the device.)
6. If the device is explicitly permitted by a device access rule, it's granted full access.
7. If the device state hasn't changed during any of the preceding steps, then the default access state that's set in the EAS organizational settings is applied.

Speaking of organizational settings, the most important setting is the one that controls what happens when a device connects and doesn't have any ABQ settings in place. The default is to allow such devices to connect; that is, if the device isn't explicitly blocked, then it's allowed. If you want to change this behavior, you can do so by using the settings that Figure 2 shows.

Figure 2
Changing Default Device Behavior

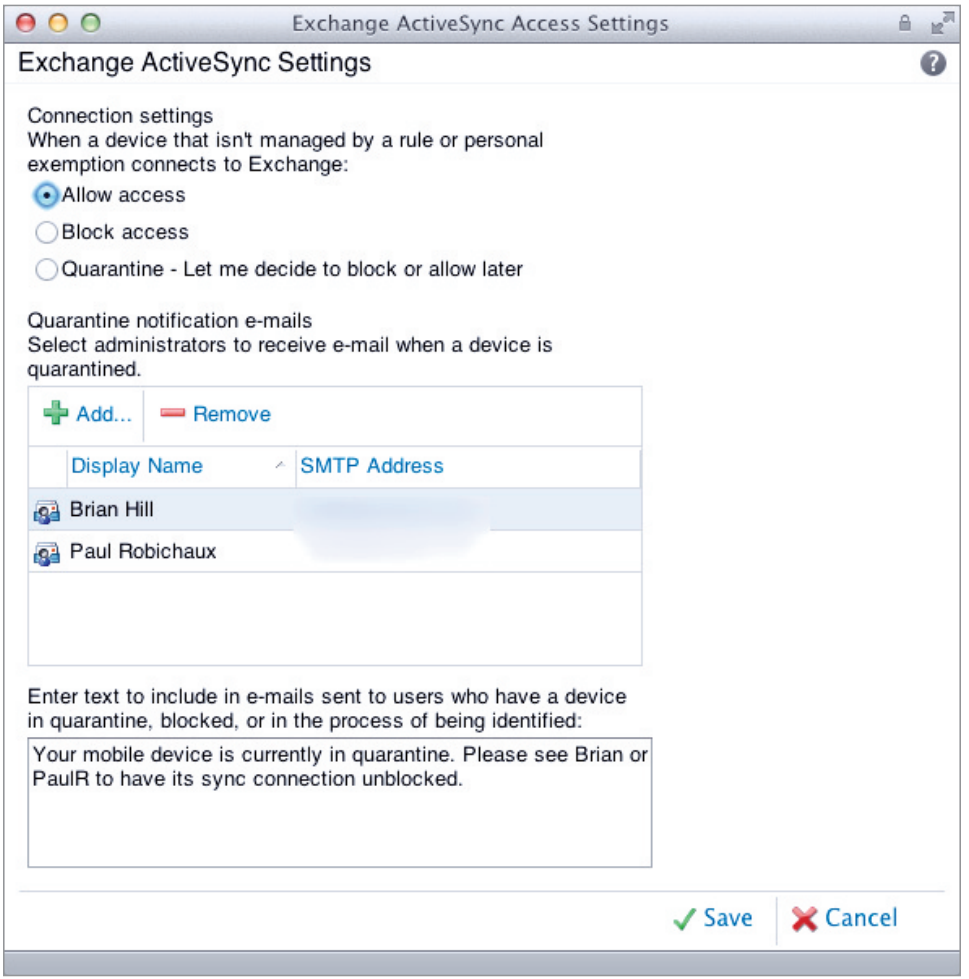

You'll see this dialog box when you click Edit on the ActiveSync Access tab of the Phone & Voice slab in ECP. The *Allow access*, *Block access*, and *Quarantine – Let me decide to block or allow later* radio buttons control what happens to devices when sync permissions aren't otherwise specified. The *Quarantine notification e-mails* list shows which administrators will receive email messages when a device is quarantined. The text field at the bottom of the dialog box allows you to add text that will appear in the default quarantine message that users receive when their devices are quarantined.

Blocking, Quarantining, and Allowing Devices

Now that you know how EAS ABQ is implemented, you're probably wondering how you actually set things up to control which devices can connect. You need to worry about only a few scenarios:

- When you want to block everyone from using all devices, with a few well-specified exceptions, set the organizational EAS policy to block devices, then add personal exceptions for the users and devices that you want to let synchronize.
- When you want to know before any user syncs any device, use the quarantine policy, which is made expressly for this kind of situation. Enable quarantine on the organizational EAS policy and every user who connects a device will be quarantined unless and until you release them.
- When you want to allow some (or all) users to use specified devices only, set the organizational EAS policy to block or quarantine devices, then create a device access rule to allow the devices and families that you want. If there are users you want to exclude, simply turn off EAS on their mailboxes.
- When you want to allow some (or all) users to sync with any kind of devices they want but require other users to get permission, set the organizational EAS permission to quarantine. When a previously unsynchronized device tries to join the party, it will be quarantined unless a personal exemption exists.

The Future of ABQ

The enterprise computing world is being swept by "bring your own device" mania. Most companies don't want to provide—or pay for—their employees' mobile devices, so they have instead refocused on how to control which devices are allowed to access company assets such as Exchange servers. In that light, the existing ABQ framework is quite useful: It provides tools for allowing or blocking device access based on who the device belongs to or the type of device.

The built-in Windows Mail application in Windows 8 uses EAS and can apply EAS policies. It's too early to tell whether EAS support is a trend, but a key irritation with Outlook Anywhere is that you can't easily control which machines users connect from. Having an EAS-based mechanism for applying security policies to computers running Outlook would be a major improvement, as would better Exchange tools for specifying criteria for allowing, blocking, or quarantining devices.

Despite the allure of these potential improvements, the existing Exchange 2010 EAS ABQ features are quite useful on their own. These features provide a solid foundation for controlling which devices users can use to sync their mailbox data. ■

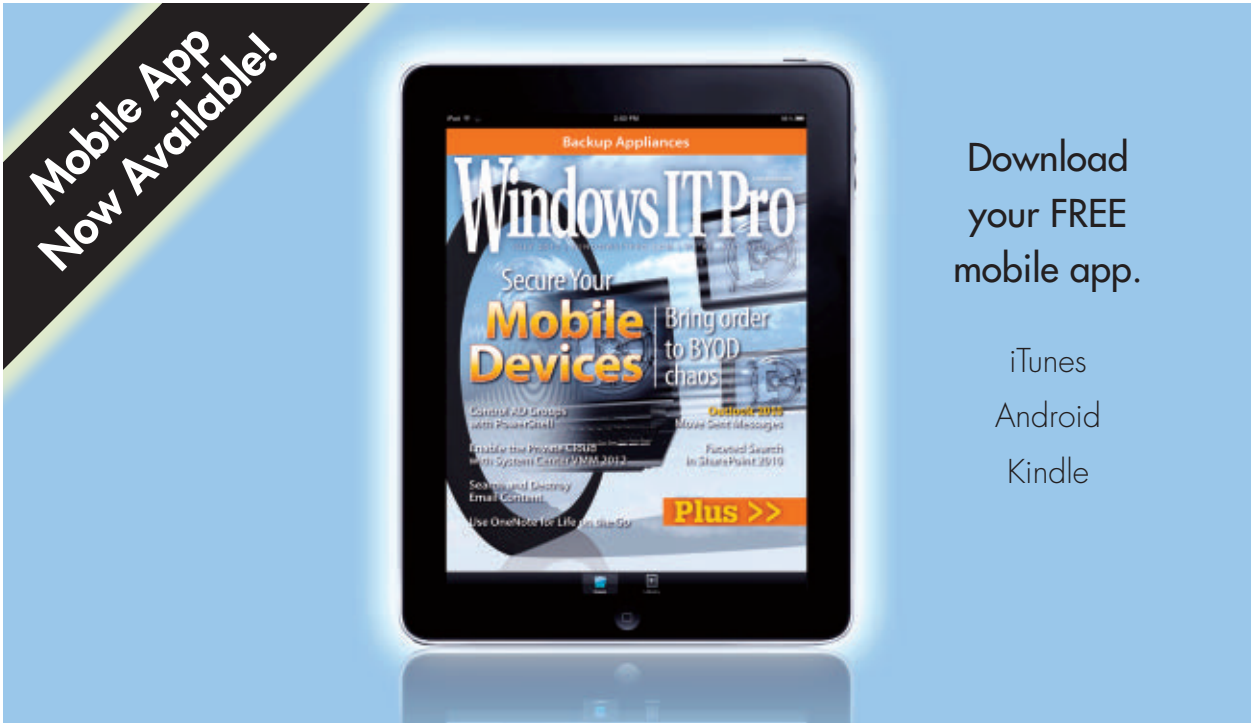

Using the Confidentiality Bit to Hide Data in Active Directory

Configure which attributes can replicate to an RODC

Microsoft Active Directory (AD) has decent capabilities for setting permissions on objects. You can use these permissions to allow delegated administration of users, groups, or computers to any security principal. In this way, many daily operations don't need to be performed by domain administrators. But when it comes to making specific data visible to only those users who need to see it—either because typical users shouldn't see the objects or because the data is truly confidential—the default AD permissions can make the task rather complex.

This four-part series has discussed AD data-hiding options. These options can be based on using normal AD permissions, a special AD permission feature called List Object mode (or List Mode), or a more advanced option, the confidentiality bit (a lesser-known option introduced a few years back in Windows Server 2003 Service Pack 1—SP1). With respect to setting permissions on AD data, there have been only minor enhancements in Windows Server 2008 R2 and Windows Server 2008. I'll describe those enhancements in this article. (See the Learning Path for a list of the previous articles in the series.) I'll also describe use of the confidentiality bit, as well as how it relates to configuring which attributes are replicated to a read-only domain controller (RODC).

Guido Grillenmeier

is a chief engineer within the Enterprise Services Group at HP. He is a Microsoft Directory Services MVP, a Microsoft Certified Architect, and the coauthor of *Microsoft Windows Security Fundamentals* (Digital Press).

Email

Using the Confidentiality Bit

The last of the various options to hide confidential data in AD is the confidentiality bit. This feature was added to AD specifically to support another Windows Server 2003 SP1 security-related feature:

the Credential Roaming feature, officially called the Digital Identity Management Service (DIMS).

DIMS enhances the options that are available for storing a user's master key, which is used for encrypted data such as the Encrypting File System (EFS). This key is typically stored in the user profile, which causes various challenges when a user works on multiple machines and thus has different master keys. Although roaming profiles allow users to use the same master key on multiple machines, the profiles have their own challenges. With the introduction of DIMS, Microsoft offers a way to store the master key files in AD by extending the schema with a few extra attributes (i.e., ms-PKI-DPAPIMasterKey, ms-PKI-AccountCredentials, and ms-PKI-RoamingTimeStamp) for the userClass schema object. (The schema isn't extended automatically to make these attributes available; you must add them separately.)

This article won't discuss DIMS in great detail, but from the short introduction you can see that the new DIMS feature stores very sensitive data—the master key files—on a user object.

As I explain in the previous articles in this series, each user has Read permission to all the attributes of his or her user object; the well-known SELF security principal is granted the Read All Properties permission to a user object. Also, many companies grant a special group the Read All Properties permission to a whole tree of objects (if not the whole domain or forest). For example, you can use this approach to allow certain service accounts to read any data from AD, without requiring domain admin privileges.

So how do you prevent access to sensitive data for accounts that have been granted Read access to certain attributes either directly, through property sets, or through the all-encompassing Read All Properties permission? Sensitive data might include the attributes that hold the master keys, employee social security numbers, or even employee IDs—whatever your company considers to be sensitive. This is where the confidentiality bit comes in.

The confidentiality bit was introduced in Windows Server 2003 SP1. As the name implies, it configures specific attributes in AD to be confidential; normal Read privileges are insufficient to be granted Read access on these attributes.

I'll cover how to mark an attribute as confidential and how to grant access to read a confidential attribute in a moment, but the basic idea is to do the following:

1. The new confidentiality bit is set as bit 7 (= 128 decimal) in the searchFlags property of the respective attributeSchema object in the AD schema. Add 128 to any existing value to designate the attribute as confidential.
2. To grant access to a confidential attribute to users or groups that need to read the confidential data in the attribute, you must give them the CONTROL_ACCESS permission on the attribute for the respective objects. This introduces a way for AD to impose additional security checks that control Read access to selected attributes.

When trying to use the confidentiality bit, be aware of this important limitation: Microsoft doesn't let you apply the bit to the base schema attributes. In other words, you can't leverage the confidentiality bit for 90 percent of the default attributes that come with AD.

The base schema attributes are Category 1 attributeSchema objects and can be identified by using their systemFlags attribute, as bit 4 (= 16 decimal). But not all default attributes are Category 1 attributes. The AD schema comes out of the box with 863 attributes in Windows 2000 and 1,070 in Windows Server 2003. Windows Server 2003 R2 adds another 81 attributes (mostly for Services for UNIX—SFU—and DFS Replication—DFSR). Since then, the Windows Server OS releases have added further attributes to the base schema. By now, you'll find 1,264 base schema attributes in Windows Server 2012.

To figure out which attributes are Category 1, we'll use the native LDAP query tool LDP.exe to search for all attributeSchema objects with

bit 4 enabled. You could simply dump all attributeSchema objects, along with their systemFlag property, with a filter such as (objectCategory = attributeSchema) and then perform the analysis using some other tool, such as Microsoft Excel. But it's much nicer to run an LDAP query that gives you the final result right away. Yet we can't run a query checking for a simple decimal value in the systemFlag property; an attribute can have other bits set in systemFlags as well. We need to run an LDAP query with a bitwise test, by adding the RuleOID 1.2.840.113556.1.4.803 into the search filter. There are two RuleOIDs for bitwise operations (matching rules): 1.2.840.113556.1.4.803 is an AND condition (i.e., true only if all bits of the decimal value are matched), and 1.2.840.113556.1.4.804 is an OR condition (i.e., true if any bit of the decimal value is matched). As such, our LDAP search filter would be as follows:

```
(&(objectCategory=attributeSchema)
  (systemFlags:1.2.840.113556.1.4.803:=16))
```

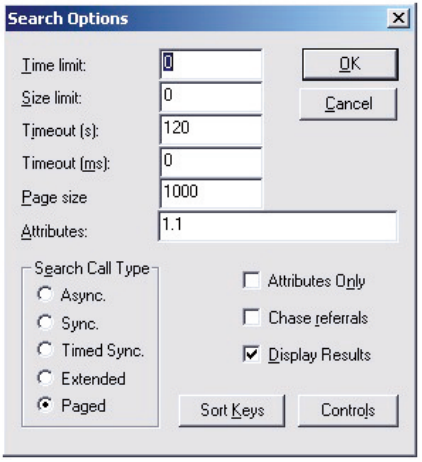

Figure 1
Options for Running a Paged Query with LDP

We expect a value above 1,000, so we need to enable a paged query in LDP. Open the Search Options dialog box, and set the Search Call Type to Paged, as shown in Figure 1. Enter the Object Identifier (OID) 1.1 in the Attributes field to tell LDP to return only distinguished names (DNs) and no attributes. Depending on the version of LDP that you use, the Attributes field might be visible only in the Search dialog box, where you can also enter the value 1.1 for the same result.

The results of the paged query against the Schema naming context (NC) of a Windows Server 2003 AD, using the filter shown in Figure 1, returns 1,007 attributes that belong to Category 1, as Figure 2

Figure 2: Filtering for Category 1 Attributes

```
***Searching...
ldap_search_init_page(ld, "CN=Schema,CN=Configuration,DC=root,DC=net", 2,
  "(&(objectCategory=attributeSchema)(systemFlags:1.2.840.113556.1.4.803:=16))",
  attrList,  0, svrCtrls, ClntCtrls, 0, 0 ,SortKeys)
0x0 = ldap_get_next_page_s(ld, hPage, 1000, &timeout, 0, 0xC05750);
Result <0>:
Matched DNs:
Getting 1000 entries:
>> Dn: CN=Account-Expires,CN=Schema,CN=Configuration,DC=...
>> Dn: CN=...
>> Dn: CN=...
>> Dn: CN=Well-Known-Objects,CN=Schema,CN=Configuration,DC=...
-----------
   -=>> 'Run' for more, 'Close' to abandon <<=-
0x0 = ldap_get_next_page_s(ld, hPage, 1000, &timeout, 0, 0x9549D8);
Result <0>:
Matched DNs:
Getting 7 entries:
>> Dn: CN=When-Changed,CN=Schema,CN=Configuration,DC=...
>> Dn: CN=...
>> Dn: CN=...
>> Dn: CN=X509-Cert,CN=Schema,CN=Configuration,DC=...
```

shows. (Windows Server 2003 R2 doesn't add any Category 1 attributes to AD.)

If you run the query, you'll notice that many potential confidential AD attribute candidates are actually Category 1 attributes and thus can't be used in combination with the confidentiality bit. These include the attributes in Table 1.

As such, it's advisable to slightly edit the LDAP query to show only those attributes that aren't Category 1 (i.e., don't have bit 4 of systemFlags enabled). You can do this by adding a NOT expression to the filter:

```
(&(objectCategory=attributeSchema)
  (!(systemFlags:1.2.840.113556.1.4.803:=16)))
```

The result is 63 attributes for Windows Server 2003 and 144 attributes for Windows Server 2003 R2, all of which can potentially be used with the confidentiality bit. With Windows Server 2012, you'd even get 164 attributes. To put this into perspective, the attributes in Table 2 are only those attributes in a Windows Server 2003 AD forest that can be used with the user class object, bringing the number down to 25 attributes.

Table 1: Excerpt of List Attributes that Can't Be Used with the Confidentiality Bit

CN	lDAPDisplayName
Comment	info
Description	description
Employee ID	employeeID
Phone-Home-Primary	homePhone
Phone-Home-Other	otherHomePhone

Table 2: Default User Attributes that Can Be Used with the Confidentiality Bit

CN	lDAPDisplayName
Address-Home	homePostalAddress
attributeCertificateAttribute	attributeCertificateAttribute
audio	audio
carLicense	carLicense
departmentNumber	departmentNumber
Employee-Number	employeeNumber
Employee-Type	employeeType

Table 2: *continued*

CN	lDAPDisplayName
houseIdentifier	houseIdentifier
jpegPhoto	jpegPhoto
labeledURI	labeledURI
ms-DS-Object-Reference-BL	msDS-ObjectReferenceBL
ms-Exch-Assistant-Name	msExchAssistantName
ms-Exch-House-Identifier	msExchHouseIdentifier
ms-Exch-LabeledURI	msExchLabeledURI
Network-Address	networkAddress
Other-Mailbox	otherMailbox
photo	photo
preferredLanguage	preferredLanguage
Registered-Address	registeredAddress
roomNumber	roomNumber
secretary	secretary
Text-Encoded-OR-Address	textEncodedORAddress
userPKCS12	userPKCS12
User-SMIME-Certificate	userSMIMECertificate
x500uniqueIdentifier	x500uniqueIdentifier

There isn't much sense in questioning why Microsoft specifically chose these attributes to be base schema attributes. Basically, Microsoft wanted to limit the use of the confidential data feature to custom AD extensions, such as an attribute containing the Social Security number of a user, which isn't part of the default schema.

Nevertheless, some companies might want to store the employeeID in AD. They would like to store this data in an attribute that can be changed to a confidential attribute so that only authorized users

can read and edit it. If that's the case, they can choose not to store the data in the attribute called employeeID, but instead to use the employeeNumber attribute. The latter isn't a Category 1 attribute. As such, it can be configured as a confidential attribute. Because its searchFlag attribute is empty, writing the value 128 into the searchFlags property of the employeeNumber attribute in the schema, as Figure 3 shows, is sufficient. If the searchFlags property isn't empty, then you need to add the number. The searchFlags attribute defines various other options for an attribute, such as whether it's indexed (bit 1) or remains in the tombstone object at deletion (bit 3). For more details about searchFlags, check out the Microsoft article "Search-Flags Attribute (Windows)."

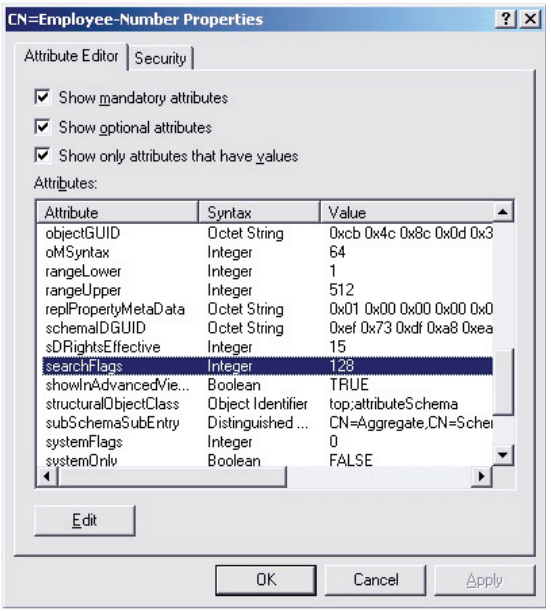

Figure 3 Setting the Confidentiality Bit for employeeNumber with ADSIedit

If you try to set the confidentiality bit on a Category 1 attribute (base schema), you'll receive a misleading error message, as shown in Figure 4. Always remember that base schema attributes can't be made confidential.

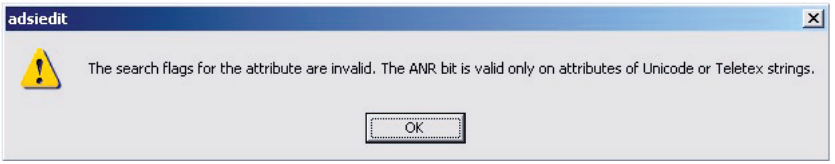

Figure 4 Error when Trying to Set the Confidentiality Bit for a Base Schema Attribute

As soon as this flag is set (and the schema-cache is updated), a populated employeeNumber attribute is no longer be visible to a user

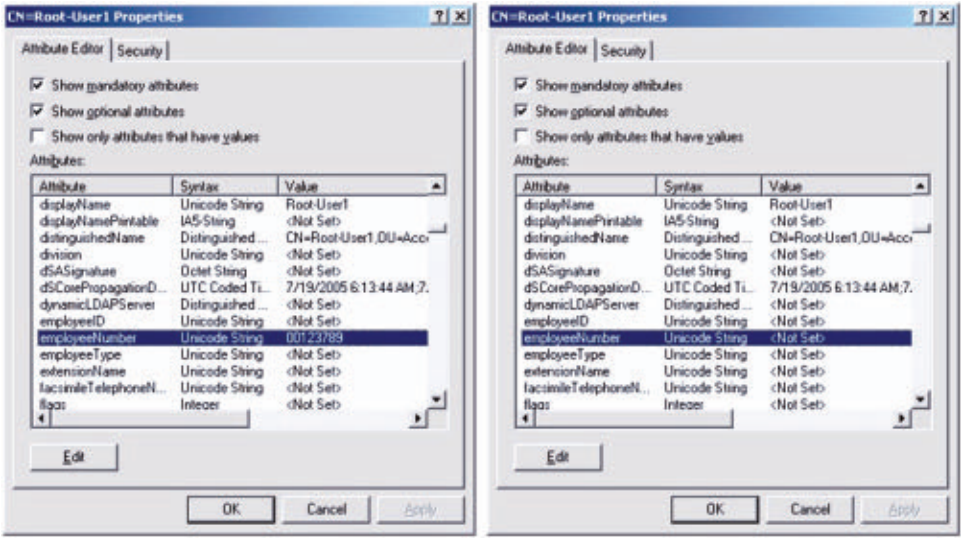

Figure 5
Contents of employeeNumber as SELF Before (Left) and After (Right) Activating the Confidentiality Bit

who is granted the Read All Attributes via the SELF security principal in the default permissions. Figure 5 shows the contents of the attribute before (on the left) and after (on the right) activating the confidentiality bit.

Another challenge is already ahead of us: How do we set the CONTROL_ACCESS permission for a hidden attribute? This should actually be an easy thing, but Microsoft didn't supply any command-line tools with Windows Server 2003 SP1 that could set this access at the attribute level. However, Windows Server 2008 and later versions have an updated Dsacls version that fully supports this capability, as tested on a Windows Server 2012 DC.

The correct syntax to add the CONTROL_ACCESS permission via Dsacls is as follows:

```
DSACLS <DN of object> /G <secprin>:CA;<property>
```

When you assign CONTROL_ACCESS permissions at the property level to a user or group, you must specify the display name of the property—in our case, employeeNumber:

```
DSACLS "CN=Root-User1,OU=UserAccounts,DC=root,DC=net"
  /G root\HR-users:employeeNumber
```

Unfortunately, ever since the CONTROL_ACCESS permission was introduced, there hasn't been a really useful UI to manage or view this permission. This is still true in Windows Server 2012. But since Windows Server 2003 R2, the LDP.exe editor does include a powerful security editor that allows you to view and set the CONTROL_ACCESS flag on a specific object attribute.

Navigate to the object on which you want to change the permissions. Right-click the object and choose Advanced, Security Descriptor. LDP then pops up a dialog box, which you can use to set the options for displaying the Security Descriptor. Don't choose the Text dump, which dumps the descriptor to the output window. Using the default settings starts the new Security Editor (see Figure 6). Choose Add

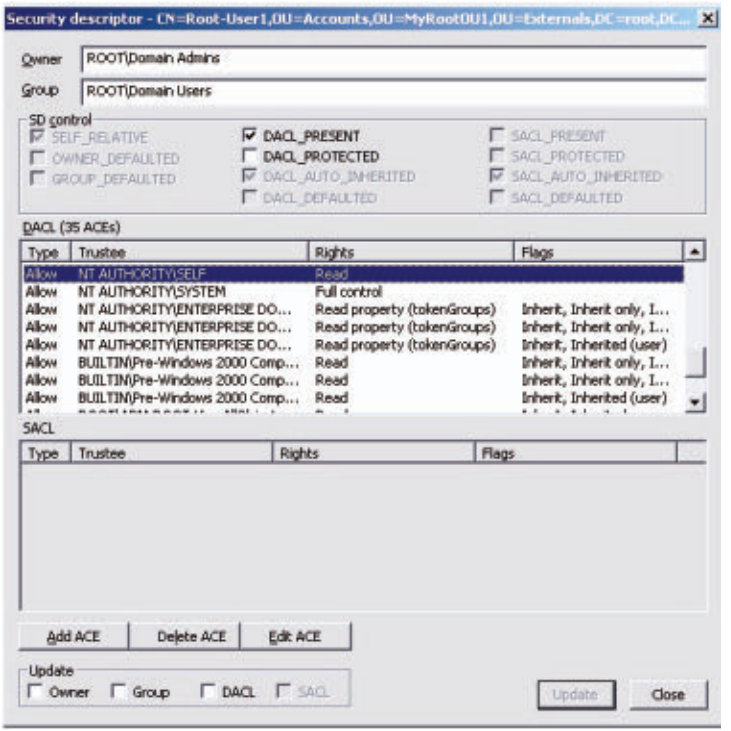

Figure 6
Security Editor from LDP.EXE in Windows Server 2003 R2

ACE to grant the control access permission for the employeeNumber attribute (see Figure 7).

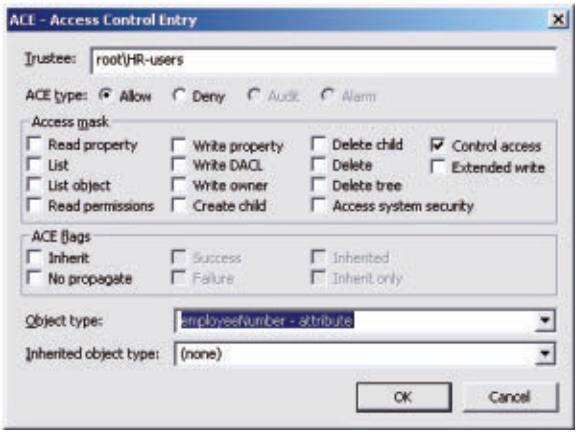

Figure 7
Setting Control Access for a Property via LDP.EXE in Windows Server 2003 R2

Managing the permissions for attributes on a list of separate objects in AD doesn't work very efficiently when using a UI. If you need to manage the permissions on confidential attributes and are still operating your AD on a version earlier than Windows Server 2008, consider setting up a tool server with a newer Windows Server OS. Doing so will allow you to use the newer Dsacls version—it's worth the effort!

RODC Filtered Attribute Set

Hiding data in AD takes yet another twist when it comes to the RODC, which was the key architectural change in AD with Windows Server 2008. Although the confidentiality bit allows us to hide data in attributes from particular users with normal Read permissions, you might want to hinder the same data from replicating to locations where you've deployed RODCs. After all, the concept of an RODC is that of a DC located in an untrusted environment.

To give you this extra security control, Microsoft added the filtered attribute set (FAS) to the RODC replication logic. This feature essentially allows you to flag attributes in the AD schema when you don't want the content of those attributes to be replicated to any RODC in the forest. As with the configuration of the confidentiality bit, you set the FAS bit via the search-flag of the attribute that you want to control; in this case, bit 9 (=512 decimal). Once set, the Read-Write DCs in the forest will no longer allow data in the respective attributes to be replicated to RODCs. Note that FAS is restricted to the same

attributes as the confidentiality bit: You can't use it for any base-schema attributes. But FAS certainly works just fine for any attributes that you have extended in your AD schema.

Stay Sensitive

You can leverage the confidentiality bit to control access to sensitive attributes. This approach allows you to hide specific attributes from users (e.g., the SELF security principal) who otherwise are granted general Read access via the default permissions on objects in AD.

You must remember a few things when considering the use of the confidentiality bit to hide attributes from general Read access in AD. As with the other data-hiding options, the confidentiality bit doesn't hinder domain or organizational unit (OU) admins from viewing the confidential attributes of an object in AD. However, it does hinder those users that are merely granted general Read permissions on the whole object or via a property that's set to read the data in an attribute that's configured to be confidential.

Marking an attribute confidential immediately changes the permissions of the respective attribute for all objects in AD. Most default attributes in AD (all of which are base schema attributes) can't be marked confidential. As with the change of Property-Sets, no re-ACLing of the objects in AD is required to activate the change of permission in an AD forest. However, to grant read access to non-admins, you need to grant the target group the CONTROL_ACCESS permission to the confidential attribute.

Although it took a while for Microsoft to properly update the Dsacls command-line tool to efficiently manage the required attribute permissions, with the latest server versions, there's no excuse for not properly configuring your AD permissions to hide confidential data. ■

Learning Path

"Hiding Data in Active Directory"

"Hiding Active Directory Objects and Attributes"

"Enabling List Object Mode in a Forest"

FAQ

Answers to Your Questions

John Savill

Q: How can I remove Microsoft Lync–specific attributes for users?

A: If in a lab or perhaps a failed production Lync deployment, you enabled users, then didn't remove the users from Lync before de-provisioning the Lync server, the users will keep their Lync attributes. If you try to add users to the new Lync server, they won't be listed because their attributes are already populated.

Use the following Windows PowerShell command to remove Lync attributes:

```
get-aduser -filter {msRTCSIP-PrimaryUserAddress
  -like "*"}|set-aduser -clear msRTCSIP-PrimaryUserAddress,
  msRTCSIP-PrimaryHomeServer,msRTCSIP-UserEnabled,
  msRTCSIP-OptionFlags,msRTCSIP-UserPolicies
```

Since it removes the Lync attributes from all users in the domain, use this only in a lab environment or when you're absolutely sure that Lync has been removed in your environment.

—John Savill

Q: How do I take a screen shot on a Windows RT device?

A: To take a screen shot on a Windows RT device, press the Windows button and the volume button at the same time. The screen will dim slightly, and in your Pictures folder you'll see that

Figure 1
Windows RT Device Screenshots Folder

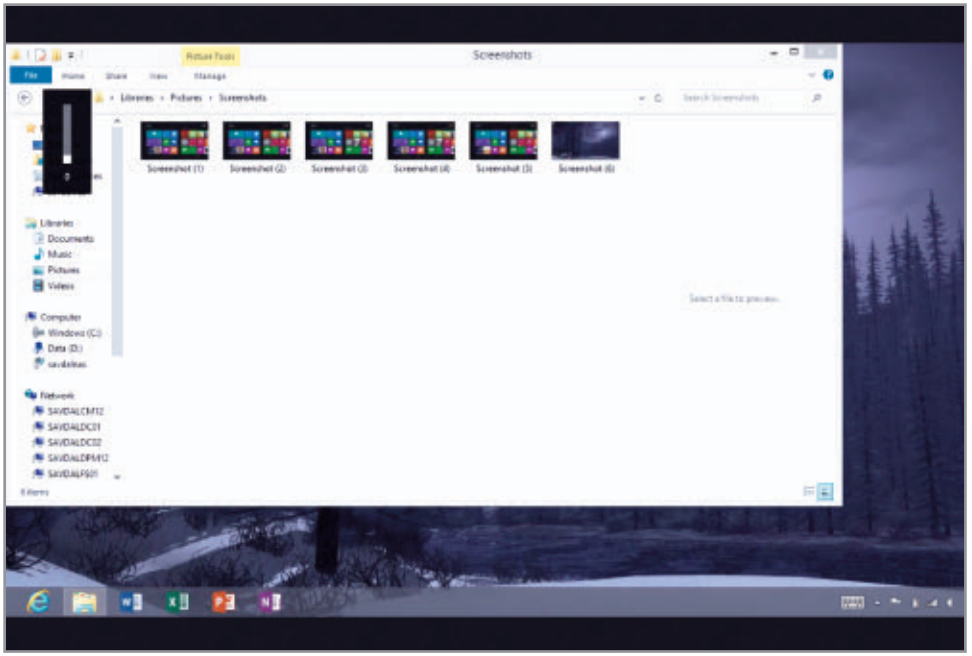

a folder has been created, labeled Screenshots. In it you'll find your screen shot, as Figure 1 shows.

—John Savill

Q: Which version of Hyper-V is best for my workloads?

A: You have, essentially, three ways to acquire Hyper-V technology (excluding the client Hyper-V on Windows 8), and all three have exactly the same capabilities and scalability. The only differences between them are the included rights to run Windows Server on guest virtual machines (VMs) on the host.

These rights are for the included Windows Server running VMs. It's not a problem to run other OSs that are licensed, and you can even run additional Windows Server VMs, but they would have to be licensed separately. Here are the versions of Hyper-V and the rights they include:

- Windows Server 2012 Standard—provides two virtual instance rights for VMs running Windows Server.
- Windows Server 2012 Datacenter—provides an unlimited number of virtual instance rights for VMs running Windows Server.
- Microsoft Hyper-V Server 2012—provides no virtual instance rights for VMs running Windows Server (which makes sense because it's free).

Note that these rights are only for VMs running Windows Server OSs. To decide which version of Hyper-V is best for you, you should look at the use case. Here are some to consider.

VDI. You're creating a virtual desktop infrastructure (VDI) deployment that will have hundreds of VMs running Windows 8 or Windows 7. The VM guest instance rights that are part of Windows Server wouldn't be usable because the VMs aren't running Windows Server. This means the best option is to use Microsoft Hyper-V Server, which is free, then license the client OSs.

Private cloud. You're deploying a private cloud with a lot of VMs running Windows Server OSs that will move between hosts using Live Migration. Use Windows Server Datacenter, which allows an unlimited number of VMs running Windows Server as part of the Datacenter license cost for full mobility of VMs between the servers.

Linux. You're deploying a Hyper-V environment to run Linux VMs. Like the desktop scenario, the guest VM instance rights for Windows servers aren't useful here because Linux has its own licensing, and customers would have to license Linux directly. Use the free Microsoft Hyper-V server.

Branch office. Your company wants to deploy a server in a branch office with two VMs running Windows Server. Use Windows Server 2012 Standard, which includes two guests running Windows Server. If in the future more VMs were required, additional standard licenses could be purchased and "stacked" on one physical server. For example, if you buy two copies of Windows Server Standard, you have the

right to run four VMs running Windows Server. It's important to note, though, that if you had two Hyper-V servers on a location running Standard, you couldn't split the two VM rights across boxes nor could you move the VMs between the servers when required. And licenses can only be moved every 90 days. Therefore, if there were two virtualization hosts with two VMs on each, but you wanted the ability to be able to move the VMs between the boxes at any time, you would need to license each host for the high-water mark of VMs (i.e., four VMs), which means two Windows Server Standard licenses would be needed for each server. Datacenter quickly becomes the most logical financial purchase as soon as you start to increase the number of VMs and need mobility of VMs between the hosts.

Basically the decision to use Microsoft Hyper-V Server versus Windows Server Standard or Windows Server Datacenter comes down to the number of VMs you want to run, with Windows Server running within them. Therefore, the virtual instance rights of the various versions matter. For non–Windows Server VMs, all three options are the same, so you should pick the most cost-effective one.

—John Savill

Q: My virtual machine has the same name on the source and target Hyper-V server. Can I rename the replica?

A: Yes, you can rename a replica virtual machine (VM). After the Hyper-V Replica replication has been enabled and is running, renaming the replica VM is fully supported. This won't cause any interruption to the replication process because the VM GUID, not the VM name, is used between the source and target. ■

—John Savill

Product News for IT Pros

Veeam Backup Ascends to the Cloud

Veeam Software announced the immediate availability of Veeam Backup Cloud Edition. Veeam Backup Cloud Edition turns every leading public storage cloud into an easy-to-use data repository for backups, providing a powerful and affordable alternative to tape and traditional offsite backup storage. The product is cloud-agnostic, with support for 15 different public storage clouds, including Windows Azure, Amazon Simple Storage Service (S3), Amazon Glacier, Rackspace, and HP Cloud (as well as support for additional clouds built on OpenStack). Key benefits of the solution include the following: Local backups can be scheduled and copied automatically to the cloud; backups to the cloud are compressed, deduplicated, and protected with up to AES 256-bit encryption; emailed reports keep IT informed about backups copied to the cloud; and the product doesn't require you to learn cloud storage APIs—you simply enter the chosen cloud's credentials into Cloud Edition. Learn more at the Veeam website.

Network Automation Launches AutoMate 9

Network Automation announced the release of its latest automation software solutions, AutoMate 9 and AutoMate BPA Server 9. The release marks the expansion of the no-code platform's enhanced offerings, including Microsoft Azure (fully enabled cloud storage automation via Azure Storage Services); Microsoft Dynamics CRM (automation of a wide range of entity data management activities, including creation, deletion, updating, query, retrieval, storage, and reporting of all customer, product, and order data); and OCR (converting multiple text format documents into stored data is now automated, covering formats such as PDF, TIFF, JPEG, BMP, and

GIF). Additional features include revamped automation development and deployment interfaces, increasing productivity, and control of production environments. For more information, visit the Network Automation website.

UltraBac Warp Restores Data from Any Point in Time

UltraBac Software announced UltraBac Warp, providing a comprehensive form of Continuous Data Protection (CDP) for backup and disaster recovery designed for end users with minimal technical experience. UltraBac Warp's CDP is distinctive in that it is image-based and has been identified as Continuous Image Protection (CIP). With CIP, all selected volumes are automatically protected as opposed to only certain files or folders. CIP uses changed block tracking, in which only changed blocks are saved, thereby protecting all files on partitions selected for backup. This process greatly reduces the amount of backup storage space used, by eliminating the need to save an entire file every time a change is made. UltraBac Warp also provides image-based bare-metal disaster recovery capability, even to dissimilar hardware, which file-based backup products cannot do. Another feature of UltraBac Warp is its automatic pruning settings. These allow for old data to be deleted as it reaches a user-specified expiration date. Visit the UltraBac website for more information.

SOTI Introduces Pocket Controller for Android

SOTI announced the availability of Pocket Controller for Samsung Android devices, a productivity tool that enables full remote control, as well as a variety of other useful features, for smartphone and tablet devices. Pocket Controller leverages SOTI's remote control technology, allowing users to remotely view and control their Samsung Android devices over Wi-Fi or Bluetooth. The product includes features such as full remote control, screen capture, video recording, file transfer, printing, and more. If you connect a user's desktop computer to a projector or TV, the product can also be used for demonstrations, presentations,

and remote training. Pocket Controller for Android is currently available for download within the Google Play Store. For more information, visit the SOTI website.

IS Decisions Offers Microsoft-Certified Windows 8/ Windows Server 2012 File Auditing

IS Decisions announced the release of FileAudit 4, a file auditing, archiving, and reporting solution. The highlight of FileAudit 4 is its intuitive, touch-ready UI composed of "live tiles" that link to the solution's features and update in real time. It was designed to make file auditing faster, smarter, and more efficient—regardless of whether users are working on PCs, laptops, or tablets. FileAudit 4 is also the first and only file auditing solution Microsoft-certified for compliance with Windows 8 and Windows Server 2012. Increasing mobility and remote working trends can make it challenging for IT to protect sensitive data and ensure compliance with industry regulations such as SOX, FISMA, and HIPAA. Manually monitoring and auditing file access (and access attempts) across Windows servers is time-consuming and overwhelming. With a simple, agentless deployment, IT managers can quickly install FileAudit and instantly protect all file servers in their Windows environment without intrusion or the need for deployment on individual servers. Users are up and running with FileAudit and monitoring, archiving, and reporting on file access in less than three minutes. For more information, see the IS Decisions website.

Napatech Asks You to ReThink Mobile Network Analysis

Napatech released the latest brief in its "Time to ReThink" series of insights into crucial issues facing network analysis OEM vendors. In "Time to ReThink Mobile Network Analysis," Napatech argues that there is a growing need for a new class of network analysis solutions capable of keeping up with the rapid growth in mobile data traffic. "Carriers are under pressure to generate more revenue. They now recognize that providing higher quality and innovative services is the way

forward," stated Erik Norup, President and CMO, Napatech. "But to do this, you need detailed network and application visibility at key points in the network. Together with our partners, Napatech now has the capability to help Telecom Equipment Manufacturers (TEMs) fulfill that need." Napatech's products provide the tools for real-time visibility in mobile networks. New tunneling protocol support extracts information from GPRS Tunneling Protocol (GTP) and IP-in-IP tunnels to allow efficient analysis of apps and services used on mobile networks. Napatech adapters can also perform flow identification and intelligent flow distribution to multiple server CPU cores based on the contents of the tunnel, rather than on the tunnel itself. This enables TEMs and their customers to accelerate their mobile data analysis application. For more information, see the Napatech website.

Storage Made Easy Adds Support for Amazon Glacier

Storage Made Easy announced full support for Amazon Glacier, a low-cost, fully redundant data archiving storage infrastructure. Storage Made Easy's customers can now move infrequently used data from more than 35 clouds to take advantage of Amazon Glacier's extremely low price-per-gigabyte monthly cost. With Amazon Glacier, customers can reliably store large or small amounts of data for as little as $0.01 per gigabyte per month, a significant savings compared with on-premises solutions. Storage Made Easy provides a secure cloud collaboration and file-sharing platform that unifies data from public and private cloud providers while providing full audit and governance of all known files. Customers can continue to leverage the existing public or private cloud data that they are using with Storage Made Easy's Cloud platform and use Amazon Glacier for infrequently used data. For more information, see the Storage Made Easy website. ■

WaterField Designs' Ultimate Tablet SleeveCase Offers Style and Beauty

Product of the Month

Jason Bovberg

Email

Twitter

Website

As soon as I brought my new Microsoft Surface Pro home, I knew I'd need a protective case for it. The Surface is a heavy, glossy, slippery device, and I don't have to tell you that it's rather pricey. And a typical iPad-like tablet bag wasn't going to work. The Surface Pro has a distinct, widescreen form factor; it would just slide back and forth inside the more square iPad case, or any generic tablet case. I needed something that would snugly enclose the Surface with custom design dimensions. So as I began playing with my new toy, I kept an eye on the Twitterverse for the first mention of a Surface-specific carrying case.

It was WaterField Designs that was first on the scene, and early photos of the case looked extremely promising. WaterField offers three Ultimate SleeveCases designed for the Microsoft Surface. The first, for the Surface RT, has a shallower depth to accommodate the thinner device. One size fits the Surface RT all by itself ("naked") or the Surface RT with the thin Touch Cover attached. The second holds the Surface Pro, either naked or with the Touch Cover, and the third also holds the Surface Pro, with the thicker Type Cover. I really appreciate the fact that WaterField designed its Surface cases with several scenarios in mind.

I tested the Surface Pro/Type Cover variety, and I should mention that this case comfortably holds the Surface with either of the covers,

just not both at the same time. The smooth way the device slides securely into this case is extremely satisfying.

This is a sturdy, beautiful case. Its innards are fashioned with cushions of high-grade, impact-resistant neoprene (which functions as a screen cleaner), and the exterior is a ballistic nylon shell. But any tablet case needs a little style, right? For that, the Ultimate Tablet SleeveCase boasts a lead indium or brown leather add-on trim style (for $6). The opening flap has a standard nylon handle embroidered with the company logo, looking for all the world like a stylish "WTF!" shouted across the room. (I mean this in the best way possible.) There's a large, tight pocket across the back of the case—ideal for both the Surface Pen and your microfiber cleaning cloth. You can also store your power cord in this pocket, but it is a tight fit and makes the case bulky and unwieldy. One feature that confused me momentarily was the small nylon strap at the very bottom of the case, until I realized that it's designed to hold on to as you extricate your snugly encased device. I imagine the Surface would be tough to pull out of the Ultimate SleeveCase without that tiny grab-strap! It's a smart addition, particularly for this case. You can also add a shoulder strap ($5 to $22).

WaterField touts the Ultimate SleeveCase's TSA checkpoint friendliness, and I can vouch for it. I recently took a trip through Denver International Airport and Orange County's John Wayne Airport, and at neither airport did I need to remove my tablet from the case.

I should end with the admission that, like the Surface Pro itself, the case is priced on the higher end of the spectrum at about $60 (plus add-ons). But the Ultimate SleeveCase is worth its price tag. Custom-fitted and sturdy, it will give you peace of mind about an already hefty investment. And if the glances this case got at the airport are any indication, you can't do any better in the style department. Plus, there's just something very cool about that tag on the back that reads *Made in San Francisco*. ■

Send us your funny screenshots, oddball product news, and hilarious end-user stories. If we use your submission, you'll receive a *Windows IT Pro* Rubik's Cube.

Submit

Directory of Services

Search our network of sites dedicated to hands-on technical information for IT professionals.
www.windowsitpro.com

Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
www.windowsitpro.com/go/forums

News

Check out the current news and information about Microsoft Windows technologies.
www.windowsitpro.com/go/news

EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

- Cloud & Virtualization UPDATE
- *Dev Pro* UPDATE
- Exchange & Outlook UPDATE
- Security UPDATE
- *SharePoint Pro* UPDATE
- *SQL Server Pro* UPDATE
- *Windows IT Pro* UPDATE
- WinInfo Daily UPDATE

RELATED PRODUCTS

Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the web. Includes FREE access to eBooks and archived eLearning events plus a subscription to either *Windows IT Pro* or *SQL Server Pro*.
www.windowsitpro.com/go/vipsub

SQL Server Pro

Explore the hottest new features of SQL Server, and discover practical tips and tools.
www.sqlmag.com

Dev Pro

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
www.devproconnections.com

SharePoint Pro

Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and web seminars mentored by a community of peers and professionals.
www.sharepointpromag.com

Advertiser Directory

1&1 Internet 1

Windows IT Pro 2, 60

Vendor Directory

Apple 7, 8, 52, 53

Devolutions 16

Google 10, 52

IntelliAdmin 17

IS Decisions 79

ManageEngine 18

Napatech 79, 80

NeoSmart Technologies 17

Network Automation 77, 78

Salesforce.com 21

SolarWinds 17

SOTI 78, 79

Spiceworks 18

Storage Made Easy 80

UltraBac Software 78

Veeam Software 77

Viewfinity 17

WaterField Designs 81, 82

Wireshark 18

Windows IT Pro